


Internet Strategies for Technology Leaders

new architect

formerly **webtechniques**



Tech Leaders

Sixteen technology architects share their insight on emerging technologies, the state of the Internet, and the innovation process. p.18

Security Scan

Investigating biometrics: Yesterday's hi-tech spy gear could be today's answer to ineffective password-based security. p.14

Rest Easy

If autonomic, self-healing servers become a reality, you may never have to worry about an infrastructure meltdown again. p.34

\$4.95US \$5.95CAN



06>

JUNE 2002
\$4.95 USA (Canada)
Volume 7 Issue 06
Display until June 21

Sends other UNIX

“My PowerBook G4 is now running every major UNIX app that we had on our Suns, AlphaStations and SGIs – and running them faster.”

—Mark S. Cohen, Ph.D., Professor,
UCLA Brain Mapping Center

“This new OS has accomplished in a short period of time what others have struggled to do for years: bring a compelling, widely accepted GUI (called Aqua) to UNIX.”

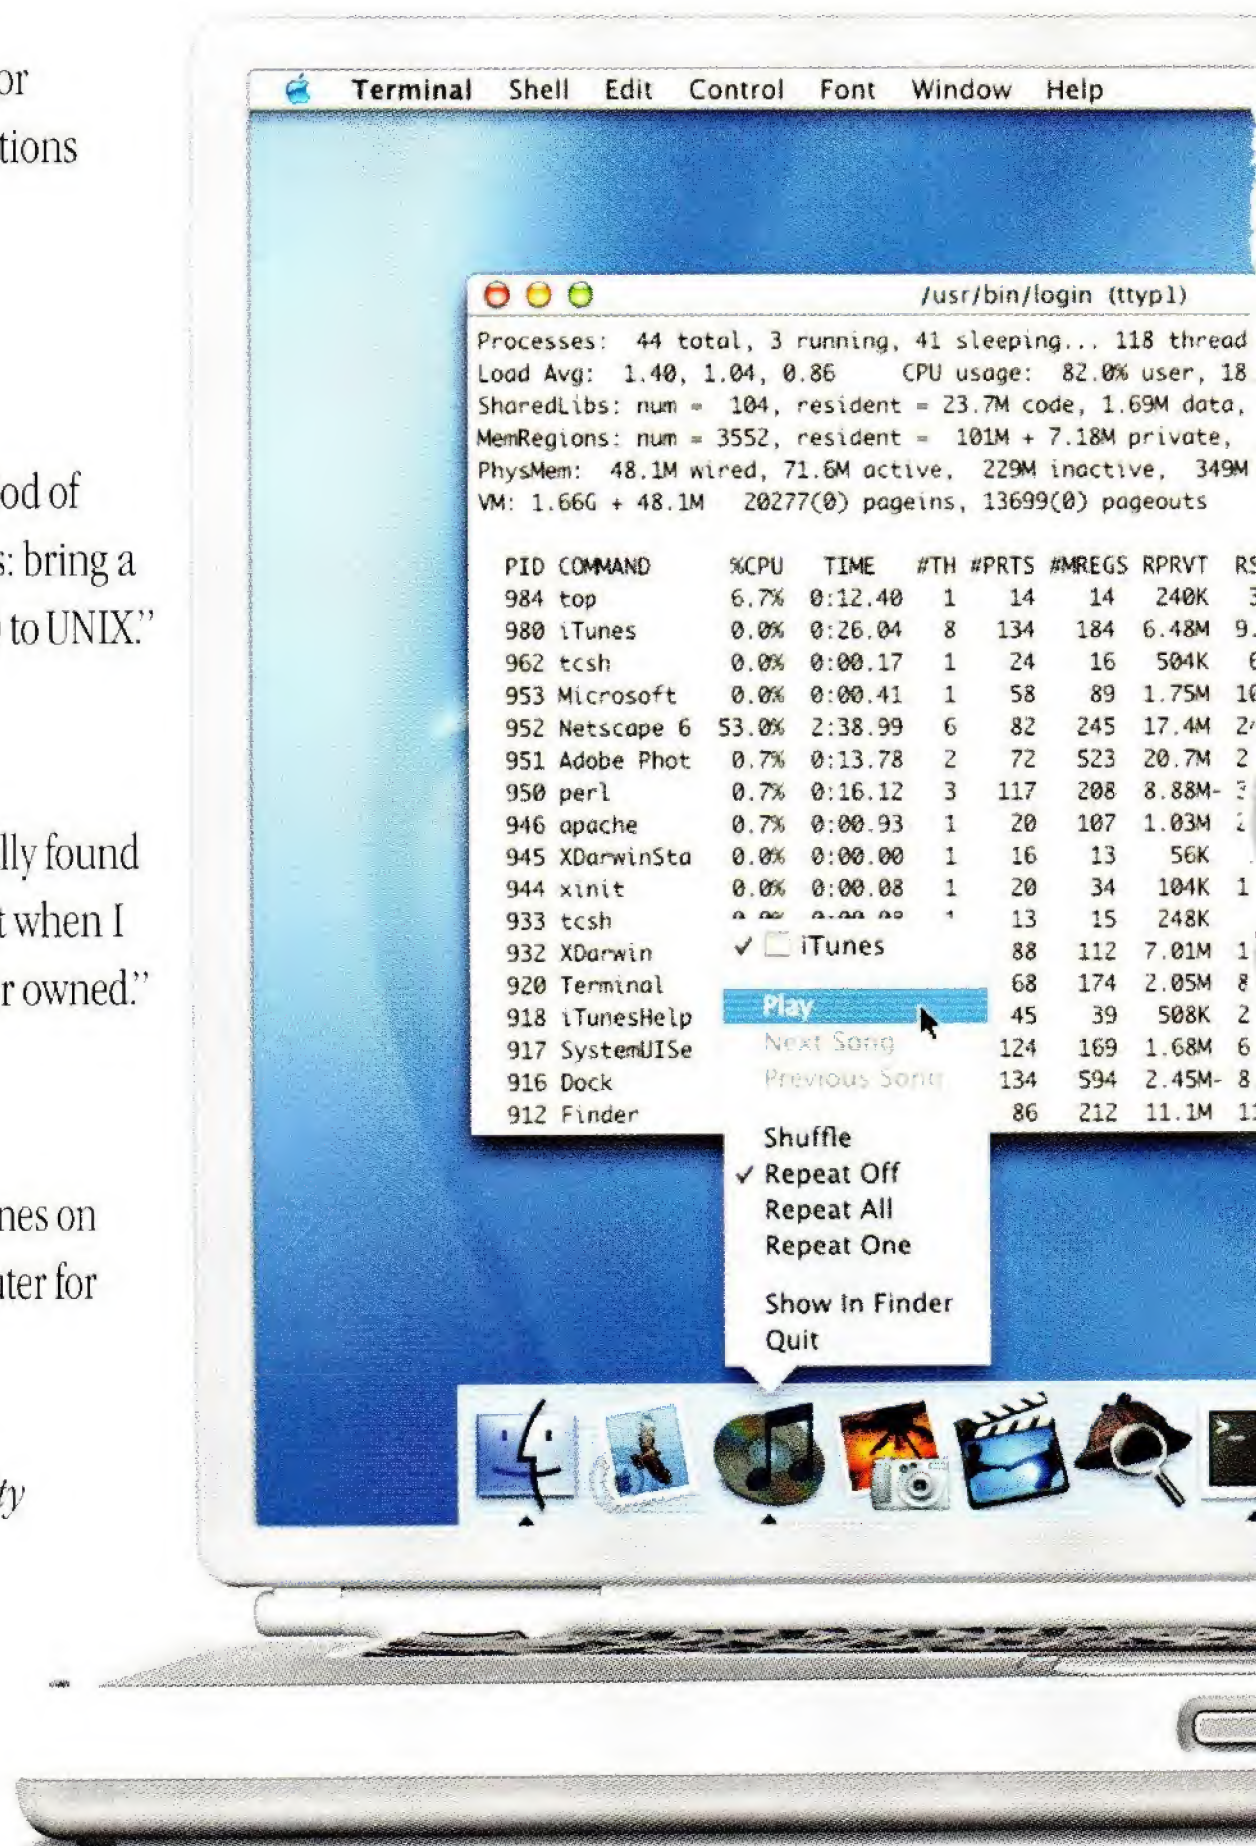
—Tim O'Reilly, O'Reilly Network

“After two-and-a-half years of Linux, I've finally found joy in a UNIX operating system. And I found it when I purchased a Macintosh – the first one I've ever owned.”

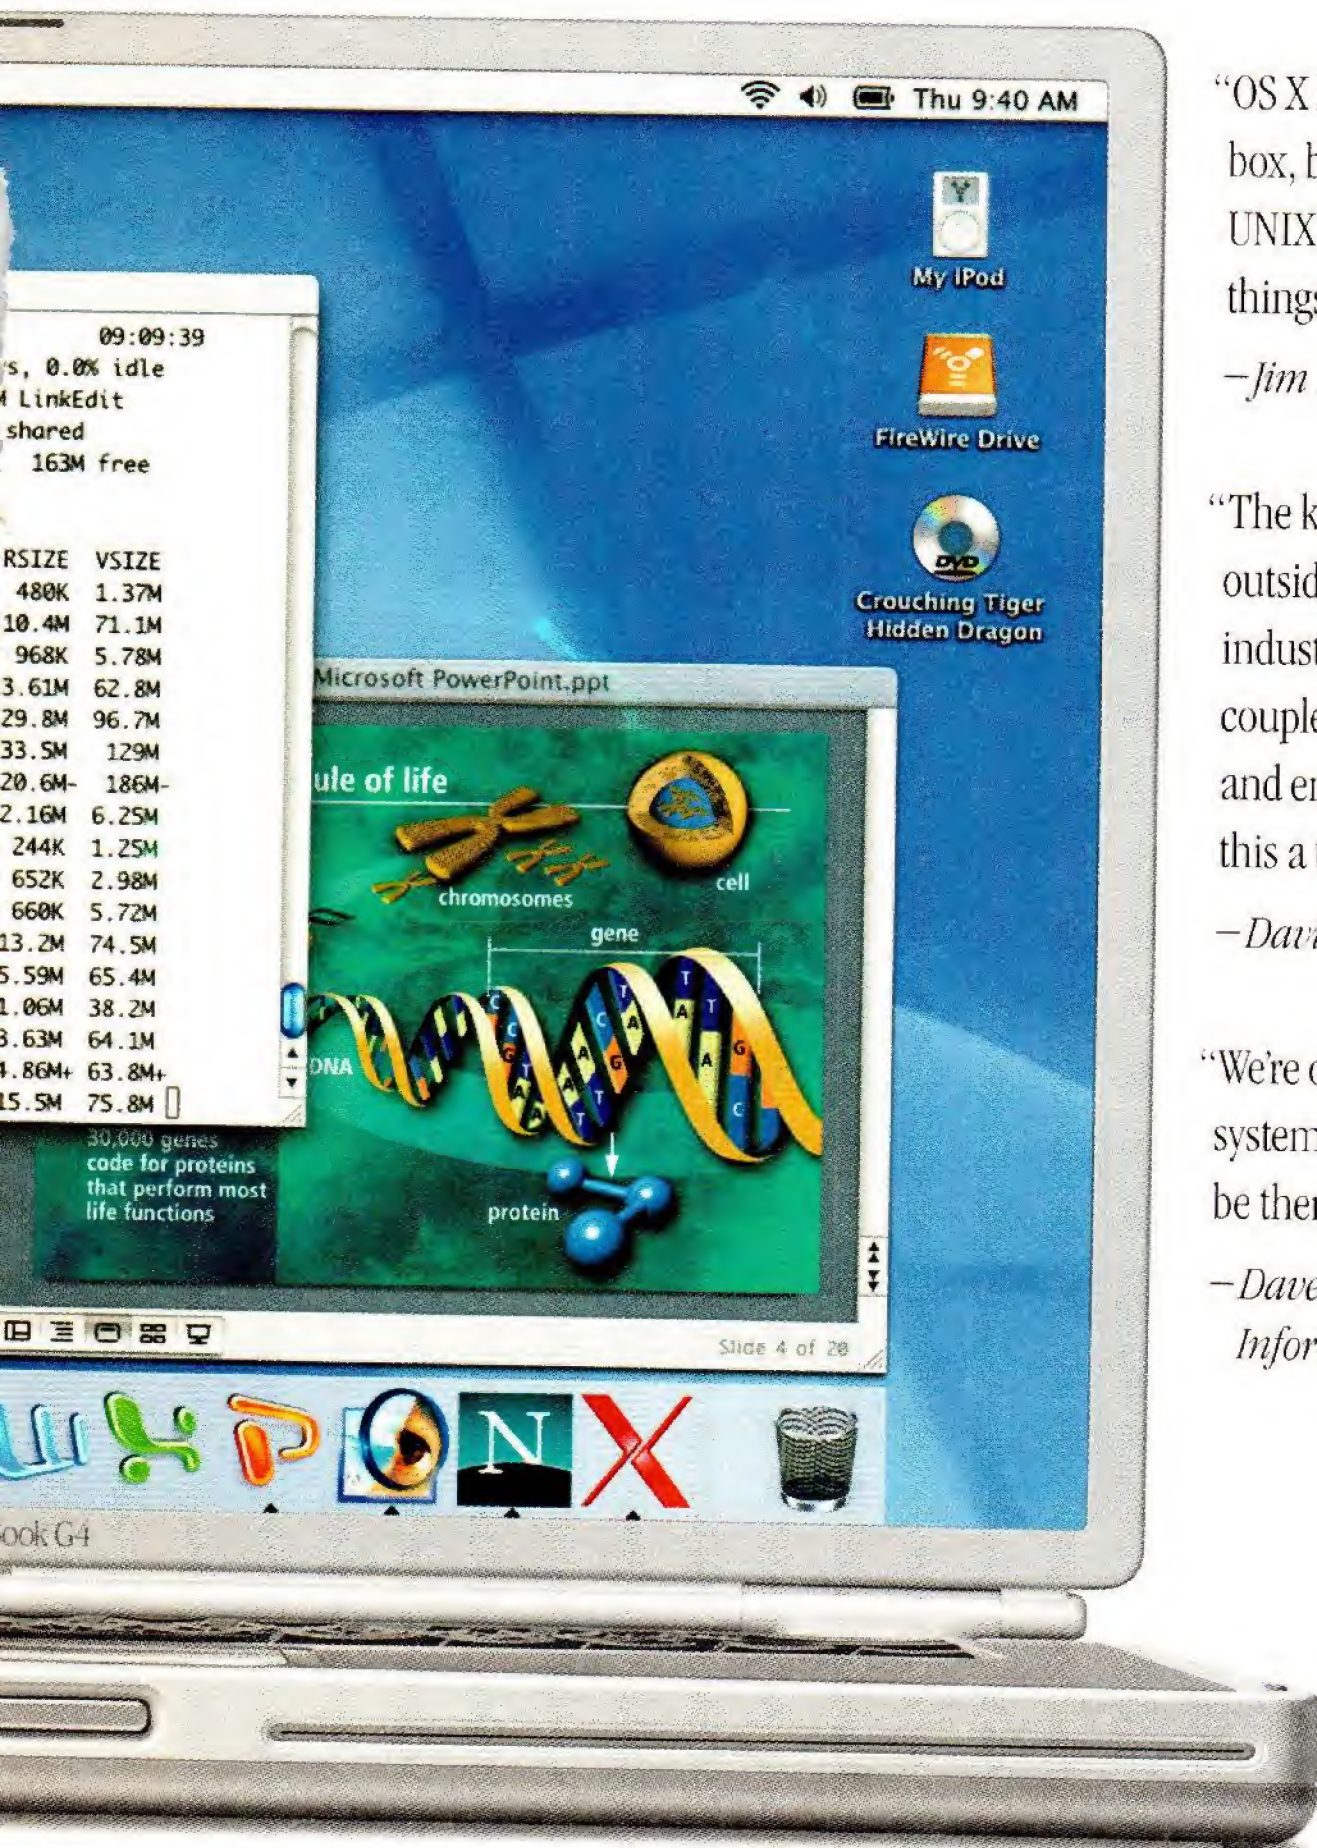
—John Hummel, Jr., The Gamers' Press

“Until Mac OS X came along, I had three machines on my desk. Now with Mac OS X, I use one computer for writing, coding, research, everything.”

—Dr. Michael Cherry, Associate Professor,
Department of Genetics, Stanford University



boxes to /dev/null.



"OS X has all the power you expect from a UNIX box, but none of the hassle associated with basic UNIX system administration. You can expect things to work and they do."

—Jim Hourihan, *Tweak Films*

"The keys to OS X are that it's UNIX inside and Mac outside. This combination of stability, support for industry standards, and wide developer support coupled with Apple's Aqua graphical user interface and emphasis on graphic performance makes this a truly great operating system."

—David Coursey, *ZDNet AnchorDesk*

"We're old hardcore UNIX hackers, so a BSD-based system is mother's milk. Everything you expect to be there is there, and it works right."

—Dave Weininger, *President, Daylight Chemical Information Systems, Inc.*



ON TIME	6155	11:05
ON TIME	927	11:10
ON TIME	1411	11:20
DELAYED	2643	11:25
ON TIME	466	11:35
ON TIME	842	11:40
DELAYED	1209	11:45
ON TIME	1744	11:55



It has been clear for some time now that the future will be a world in which everything is connected to everything, seamlessly, automatically, and securely, via XML Web services. What has not been clear is how that future will be built. Introducing Microsoft® Visual Studio® .NET, the first toolset built from the ground up for XML Web services development. Just one keyword—WebMethod—

turns your existing application into a Web service, and with new Web Forms you can create Internet applications using the same “drag-and-drop, double-click to write code” technique made famous

Microsoft
Visual Studio.net

© 2002 Microsoft Corporation. All rights reserved. Microsoft, Windows, Visual Studio, Visual Studio .NET, Visual Studio .NET logo, and the .NET logo are either registered trademarks or trademarks of Microsoft.

81A

74

Introducing
Visual Studio .NET

62

31

77A

49

66

The world takes
a little step
into the future.

LONDON

COPENHAGEN

MARS

MILAN

TOKYO

MONTREAL

NEW YORK

BUENOS AIRES



by Microsoft Visual Basic®. Plus, the whole process of building Windows®, Web, and mobile applications gets faster and easier with the new Integrated Development Environment (IDE), which gives you one toolbox, one debugger, one task list, one XML designer, and more. And within the IDE you get support for over 20 languages, including Visual Basic .NET, Visual C++®.NET, and new Visual C#™.NET, so you can leverage your existing skills and code across development teams and projects.

Start building the future: visit msdn.microsoft.com/vstudio **Software for the Agile Business.**

Microsoft®

"There is no reason
for any individual
to have a computer
in their home."

1977

"Portal,
schmortal."

2002

(Pooh-pooh it now, while you still can.)

CleverPath™ Portal and Business Intelligence Solutions

Every major innovation in technology has initially been met with more than a little resistance. Portal technology is no different. But today's curiosity is well on its way to becoming a necessity. Why? Because CleverPath Portal and Business Intelligence solutions bring your applications, business processes, and data together in one location, to be shared and used like never before. That means the right information can go straight to the right people at the right time—all in a personalized way. Don't spend another day without tapping into the true value of your enterprise. Go to ca.com/cleverpath/portal and find out more now.



Computer Associates™

HELLO TOMORROW™ | WE ARE COMPUTER ASSOCIATES | THE SOFTWARE THAT MANAGES eBUSINESS™

©2002 Computer Associates International, Inc. (CAI). All trademarks, trade names, service marks and logos referenced herein belong to their respective companies.

Internet Strategies for Technology Leaders

new.architect

18 STRATEGY: COVER STORY

By Design Wisdom from the Industry

By New Architect Editors

We asked 16 technologists, thinkers, and luminaries to share how they stay ahead and which technologies they've got their eyes on. Find out how they keep innovation alive and their businesses running smoothly.

30 DEVELOPMENT

The Languages of the Semantic Web

By Uche Ogbuji

If you believe Tim Berners-Lee, the Web has more potential than is realized today. Part of that potential is held in specifications like RDF and DAML+OIL. A new Web where agents do our bidding may not be far off.

34 INFRASTRUCTURE

Server, Heal Thyself

By Jay Lyman

IBM leads the charge toward self-managing servers with a vision it calls *autonomic computing*. Is it possible that you'll never again have to worry about your pager going off at 3 a.m.?

Advanced Web Hosting and Exchange Hosting



HOSTING SERVICES:

Microsoft Exchange Server

from \$9.95/month.

LIMITED TIME OFFER:
GET 30 DAYS OF EXCHANGE
HOSTING FREE

Windows 2000 Web Servers

with unlimited
email boxes
from \$24.95/month.

FREE WEB HOSTING SET UP
WITH PROMO CODE
ARCHITECT02

Microsoft Certified
Solution Provider

Copyright Intermedia.NET, Inc. 2002. All Rights Reserved.
Microsoft Exchange Server is a registered trademark of Microsoft
Corporation. All other trademarks are property of their respective holders.



HostPilot: The Ultimate Control Panel


HostPilot™ was designed to put maximum control at your fingertips. With other hosts, you submit requests and wait. That can add hours to your day. With HostPilot, you can manage your server and make changes instantly. Requests that took days from other hosts, can be done in 30 seconds from your keyboard.

Experience the Ultimate Control Panel, first hand.
Take a test flight at www.intermedia.net



INTERMEDIA.NET

1-800-379-7729



ACCESS AT THE TIP OF YOUR TONGUE PAGE 14

critical decisions

- 14 Are Biometrics Worth the Expense?**
It's no longer the stuff of spy novels.
- 16 Do You Need MVC?**
Applying modern design patterns to Web development.

expert opinion

- 09 From the Editor: Amit Asaravala**
Privacy holes in the Whois database leave many of us feeling helpless. Registrars can help.
- 42 Legal Code: Bret A. Fausett**
The ICANN board of directors takes a vote that will damage the organization's credibility.
- 54 Crow's Nest: Lincoln D. Stein**
The Morpheus Incident: How corporate competition could stifle the Web.
- 56 Guest Editorial: Steven Champeon**
A new multimedia tool is still no match for good ol' Dynamic HTML.

case studies

- 10 Olympic Gold**
How MSNBC.com brought home the gold at the 2002 Olympics.
- 38 A Prescription for Wireless**
Mobile devices and a wireless network help Harvard Medical School students stay focused.

reviews

- 46 Caw Networks WebAvalanche 4.0**
Dedicated hardware lets you stress-test your site against real-world loads.
- 48 Macromedia Flash MX**
Develop rich client application front ends with the newest version of Flash.
- 50 Openwave's Software Development Kit and the Nokia Toolkit 3.0**
Two new WAP IDEs simplify mobile app development.
- 52 Book Review**
The Art of Innovation: Lessons in Creativity from IDEO by Tom Kelly and Jonathan Littman.

08 Letters

Readers raise concerns about DRM and the definition of "fair use," and request a comparison of JSP and ASP.Net.

New Architect
is published monthly
by CMP Media LLC
600 Harrison Street,
San Francisco, CA 94107.

Periodicals postage paid
at San Francisco, CA
and additional mailing
offices.

Postmaster: Send
address changes to
New Architect,
P. O. Box 1246,
Skokie, IL 60076-8246.
ISSN 1537-9000. GST
(Canada) #124513185.

Cover and cover story
illustrations by
Alicia Buelow.

Editorial portraits by
Peter Stemmler.

Our readers send praise, criticism, and air concerns about DRM and the interpretation of the term *fair use*.



dismayed

I've been quite disappointed since the big transformation to *New Architect*. I read your magazines because of the technical and scripting

columns. This was one of the few magazines that I could find that covered several different Web development technologies. I liked seeing code snippets. I liked seeing markup. Now that's all gone! Instead of the shining developer's resource that *Web Techniques* once was, *New Architect* has become a completely different magazine that now seems to cater to administrators and executives instead of developers.

Eric Lund

eric@coolworks.com

thanks

Just wanted to drop you a note to let you know I really enjoyed the April 2002 issue of *New Architect*. I especially liked the two articles discussing the creation of the Orbitz and BART Web sites. It's articles like these that have made *Web Techniques*, and now *New Architect*, just about the only Internet/computing magazine I look forward to receiving.

Jay Dreyer

jaydreyer@hotmail.com

.net's true beneficiary

The last page of *New Architect*'s April 2002 issue has a very good article by Barry Parr ("The Next Big Thing," p.56) about Web services. I hope all of your subscribers get to the last page to read it. Microsoft is fooling a lot of developers into thinking that they had better get on board with all of Microsoft's .Net and Web Services propaganda, or they'll be left in the dust.

I don't think the majority of developers realize that Microsoft will be the real beneficiary of .Net, not the developers or the end-user. Barry hits the nail on the head with the parts of his article about "the dream behind

Passport and Web Services" and building solutions where there is no problem.

I hope you run that article again, maybe closer to the front.

Steve Hughes

steveh@pmts.net

fair use

I enjoyed the depth and breadth of the articles about digital rights management (DRM) in the March 2002 issue of *New Architect*. The digital medium changes the way information can be stored, retrieved, and distributed. It is the creators of information who will turn publishing, in its myriad forms, upside down. They don't need publishers and distributors anymore. At some point, the creators may want the services of a publisher, but they don't need it to start. When they realize this, DRM will manifest itself as a simple method for buying and selling content digitally.

As a writer, I care very much about my copyrights. I defend them vigorously, but I also recognize that word of mouth and sharing are still the best kinds of marketing available. People share my works legally all the time when they provide links to my online writing to their friends. I'm sure some people print my articles and pass them on to others. Fair use is a wonderful marketing tool.

It puzzles me that the term *shared* is so prevalent in the dialog about DRM. Everyone knows what it means to share, and that sharing is different from stealing or conveying property illegally. Yet, the practice of sharing is too often confused with the concept of fair use. While big publishers and distributors work on ludicrous DRM methods to prevent sharing, as a writer and consumer, I will continue to take advantage of fair use. When I want the convenience of using information on my computer, I will copy a CD to my hard drive or seek out a digital copy of a written work.

I predict that the DRM methods the publishers are developing will fail because publishers are really trying to avoid fair use.

Nick Corcodilos

nick@asktheheadhunter.com

new.architect

Group Director/Publisher Bob Melk, 415-947-6425, bob@newarchitectmag.com

EDITORIAL

Editorial Director Amit Asaravala, 415-947-6413, amit@newarchitectmag.com

Managing Editor Margaret Berry, 415-947-6415, maggie@newarchitectmag.com

Art Director Richard Koscher, 415-947-6418, richard@newarchitectmag.com

Senior Editor, Technology Neil McAllister, 415-947-6717, neil@newarchitectmag.com

Associate Products Editor Mana Tominaga, 415-947-6428, mana@newarchitectmag.com

Associate Managing Editor Allison Post, 415-947-6426, allison@newarchitectmag.com

Print Designer Eva Lipson, 415-947-6404, eva@newarchitectmag.com

Webmaster H. H. Rusa Vuong, 415-947-6429, rusa@newarchitectmag.com

Editor at Large Al Williams

Contributing Editors Bret A. Fausett, Lincoln D. Stein

Reprints Stella Valdez, 916-983-6971

ADVERTISING

Advertising Director James Long, 415-947-6421, jlong@cmp.com

Strategic Accounts Director David Katch, 212-592-8245, dkatch@cmp.com

Eastern Sales Manager Chris Salem, 203-798-2215, csalem@cmp.com

Western Sales Manager Michelle Stutsman, 415-947-6416, mstutsman@cmp.com

Ad Production Coordinator Paul O'Hare, 415-947-6941

Mailing Lists Rubin Response Inc., 847-619-9800

MARKETING

Marketing Director Diana Stanich, 415-947-6411

CIRCULATION

Group Circulation Manager Stephanie Blake

Assistant Circulation Manager Robert Zermeno

Corporate Newsstand Manager Pam Santoro

SOFTWARE DEVELOPMENT GROUP

Vice President Peter Westerman

CMP MEDIA LLC

President/CEO Gary Marshall

Executive Vice President and CFO John Day

President, Business Technology Group Adam K. Marder

President, Specialized Technologies Group Regina Starr Ridley

President, Technology Solutions Group Robert Faletra

President, Electronics Group Steve Weitzner

Senior Vice President, HR and Communications Leah Landro

Senior Vice President, Global Sales and Marketing Bill Howard

Senior Vice President, Business Development Vittoria Borazio

Vice President, Manufacturing & Distribution Bill Amstutz

Vice President and General Counsel Sandra Grayson

Managing Director, CMP Worldwide Media U.S.

Emily Ospenson-Crume

Article Submissions: Consult our editorial calendar at www.newarchitectmag.com/calendar, then send proposals to editors@newarchitectmag.com.

New Architect on the Internet: Visit our Web site at www.newarchitectmag.com or send email to editors@newarchitectmag.com.

Subscriptions: U.S.: Free to qualified recipients; all others \$34.95 for 1 year. Canada and Mexico: \$44.95 per year. Other countries: \$54.95 per year. Canada Post International Publications Mail Product (Canadian Distribution) Sales Agreement No. 0956147. International orders must be prepaid in U.S. funds.

Back issues may be purchased for \$9 per copy (includes shipping and handling). For issue availability call 800-444-4881 (U.S. and Canada). All other countries, call 785-841-1631. Back-issue orders must be prepaid—please send payment to *New Architect*, 1601 W. 23rd St., Suite 200, Lawrence, KS 66046-0127.


For subscription questions, change of address, and orders, call toll-free 800-677-2452 (U.S.). For all other countries call 847-647-5928, fax 847-647-5972, or write to: *New Architect*, P.O. Box 1246, Skokie, IL 60076-8246. Please indicate *New Architect* on all correspondence.

International newsstand distributor: Worldwide Media Service Inc., 30 Montgomery St., Jersey City, NJ 07302, or call 212-332-7100.

Entire contents © 2002 by CMP Media LLC, unless otherwise noted on specific articles. All rights reserved.



Printed in the USA



Amit Asaravala is the editorial director of the New Architect Media Group and has been with the magazine since 1999. He has more than seven years of Internet development experience.

from
the
editor

Amit
Asaravala



patching the whois database Its one big flaw presents an opportunity

If you've been working in this industry for a while, chances are that you've already registered a personal domain name for yourself, a friend, or a family member. Owning your domain name gives you a stable email address and a personalized place for your Web site and other services. The downside is that, like most individual registrants, much of your personal information is now available in the public Whois directory.

The Whois directory was designed in the early 1980s as a network-accessible database of information about hosts and IP addresses. Today, all domain name registrars must update the Whois directory with registrant information. In addition to the dates on which your domain was purchased and renewed, that information includes your name, phone number, postal address, and email address—data that anyone can look up by querying the directory.

By United States standards, nothing in the directory is private information in the way that medical records are. But unlike public phone directories, you can't opt out of the Whois database. Hence, if you purchase a domain name and accurately complete the registration form, your contact information will appear in a publicly accessible place, where marketers can mine it and blast you with unsolicited mail, phone calls, and email.

ICANN, the group that requires accredited registrars to maintain records in the Whois directory, has agreed to consider changing the system to permit more privacy. Last summer, the group published a survey on the Web to learn more about users' concerns

with the directory. Although the survey closed in August 2001, ICANN still hasn't come to a definitive conclusion on what to do. Yet the preliminary findings, available at www.does-not-exist.org/whois/, show that 92 percent of survey respondents favor opt-in or stricter protections on their Whois data.

These who still argue that the Whois directory should continue to display registrants' contact information claim that the directory is necessary to help maintain intellectual property rights. In other words, the information makes it easy for copyright holders to contact (and sue) the owners of Web sites that abuse copyrighted material. The FBI, too, supports an open directory because it helps the government track down business owners who are guilty of fraud or other violations.

These may be valid reasons, but they don't outweigh the serious liabilities of a public contact database. Not only have junk mail marketers figured out ways around access limits, but the current Whois setup gives stalkers, racists, oppressive governments, and any other malicious parties a road map to their victims' doors. A dissident no longer has the freedom to report on the atrocities of his government if he's worried that that government will use his domain name to look up his address. Likewise, a woman with a restraining order against a violent ex-boyfriend isn't going to feel comfortable setting up her own domain. These serious fears only encourage people to either forego a domain name altogether, or enter false information on their registration forms (thus, making the records

useless to the FBI and the intellectual property activists anyway).

The good news in all of this is that the ICANN accreditation agreement lets third parties enter their name and contact information into the Whois directory on behalf of actual domain registrants. In other words, registrars could put their own names and addresses into the directory if their clients want them to. Registrars have an important opportunity to offer a new feature here. In a business where the only real competition comes from lowering the prices on domain renewals, registrars that offered a privacy feature would likely gain the upper hand in the market for non-commercial customers.

The Whois required fields are remnants of a decade when only a select few researchers and government agencies registered their domain names. It's high time we updated the system to be compatible with 21st century issues. ICANN may be progressing toward this, but in the meantime, marketers, stalkers, and oppressive governments aren't going to curb their activities. Now it's up to the registrars to take action and provide individuals with a safe way to register domains. ■

A stylized, handwritten signature of Amit Asaravala.

Amit Asaravala
Editorial Director
amit@newarchitectmag.com

When sports portal Quokka declared bankruptcy, and infrastructure provider Logictier backed out of a contract, the Web site of the 2002 Olympics was left hanging.

case study

Michael Hurwicz

olympic gold MSNBC.com wins big in Salt Lake City

In 1997, NBC was gearing up to cover the 2000 Summer Olympic Games in Sydney, Australia, and was making plans for the 2002 Winter Games in Salt Lake City, UT. While the company obviously had expertise with television coverage, executives quickly realized that they weren't ready to produce Web sites for sports events of Olympic magnitude. At the time, NBC hadn't yet created its interactive arm, NBCi. And MSNBC.com was primarily a news site with limited experience covering sports events. NBC officials knew that any Olympics sites they built would have to feature up-to-the-minute results and related news content, serving as portals to immense libraries of statistics and related information.

Rather than try to build the sites on its own, NBC made a deal with Quokka Sports, a popular sports portal and development firm. For each of the 2000 and 2002 Games, Quokka would produce two sites for NBC—Olympics.com and NBCOlympics.com. Olympics sponsorships are big opportunities for vendors to get noticed. So, in exchange for the rights to produce the sites, Quokka agreed to front millions of dollars in equity, absorb all production and hosting costs, and split ad revenues with NBC in a 50/50 deal.

From the start, Quokka knew that it would be hard to make money on the Sydney Games, largely because of limits on the site content. Olympics Web sites are deliberately restricted to protect television advertising revenues. For example, live streaming video of events is limited to just minutes per day, and can be shown only after the footage has been broadcast on TV. About 3.7 billion people—well over half of the Earth's population—watched some portion

of the 2000 Olympic Games on television. In contrast, an estimated 20 million people would log onto the Sydney sites—a minuscule revenue opportunity by comparison.

The Sydney sites were widely praised for their features, but not their profits. When the games were over, only 5.6 million unique visitors had logged on, according to IDG. Quokka wrote off this disappointment as practice, and hoped that the 2002 Salt Lake Games would be more profitable. However, the underwhelming performance of the 2000 Games Web sites made it hard to gain sponsors for the 2002 sites. Before the games even began, Quokka declared bankruptcy. It was April 2001, and the news came as a blow to NBC.

The second big blow came just weeks later. Logictier, the official Internet operations sponsor for the 2002 Games, decided that its business model wasn't working and backed out of its contract to provide hardware, security, facilities, bandwidth, and other infrastructure services for the sites. With less than a year left, NBC and the Salt Lake Organizing Committee (SLOC) suddenly found themselves without anyone to build and host the Web sites for the 2002 Games.

msnbc.com steps up

NBC and the SLOC began a furious search for a replacement Web development agency. During May and June, the committee met with representatives from Yahoo, CBS SportsLine, and Ignite Media (the producer of the National Football League Web site). They were unable to strike a deal with any of the companies. With the Olympics only eight months away, time was running short, and no one was eager

to take on a job with a compressed and inflexible timeframe, high visibility in case of slip-ups, millions in up-front expenses, and dubious financial prospects after that.

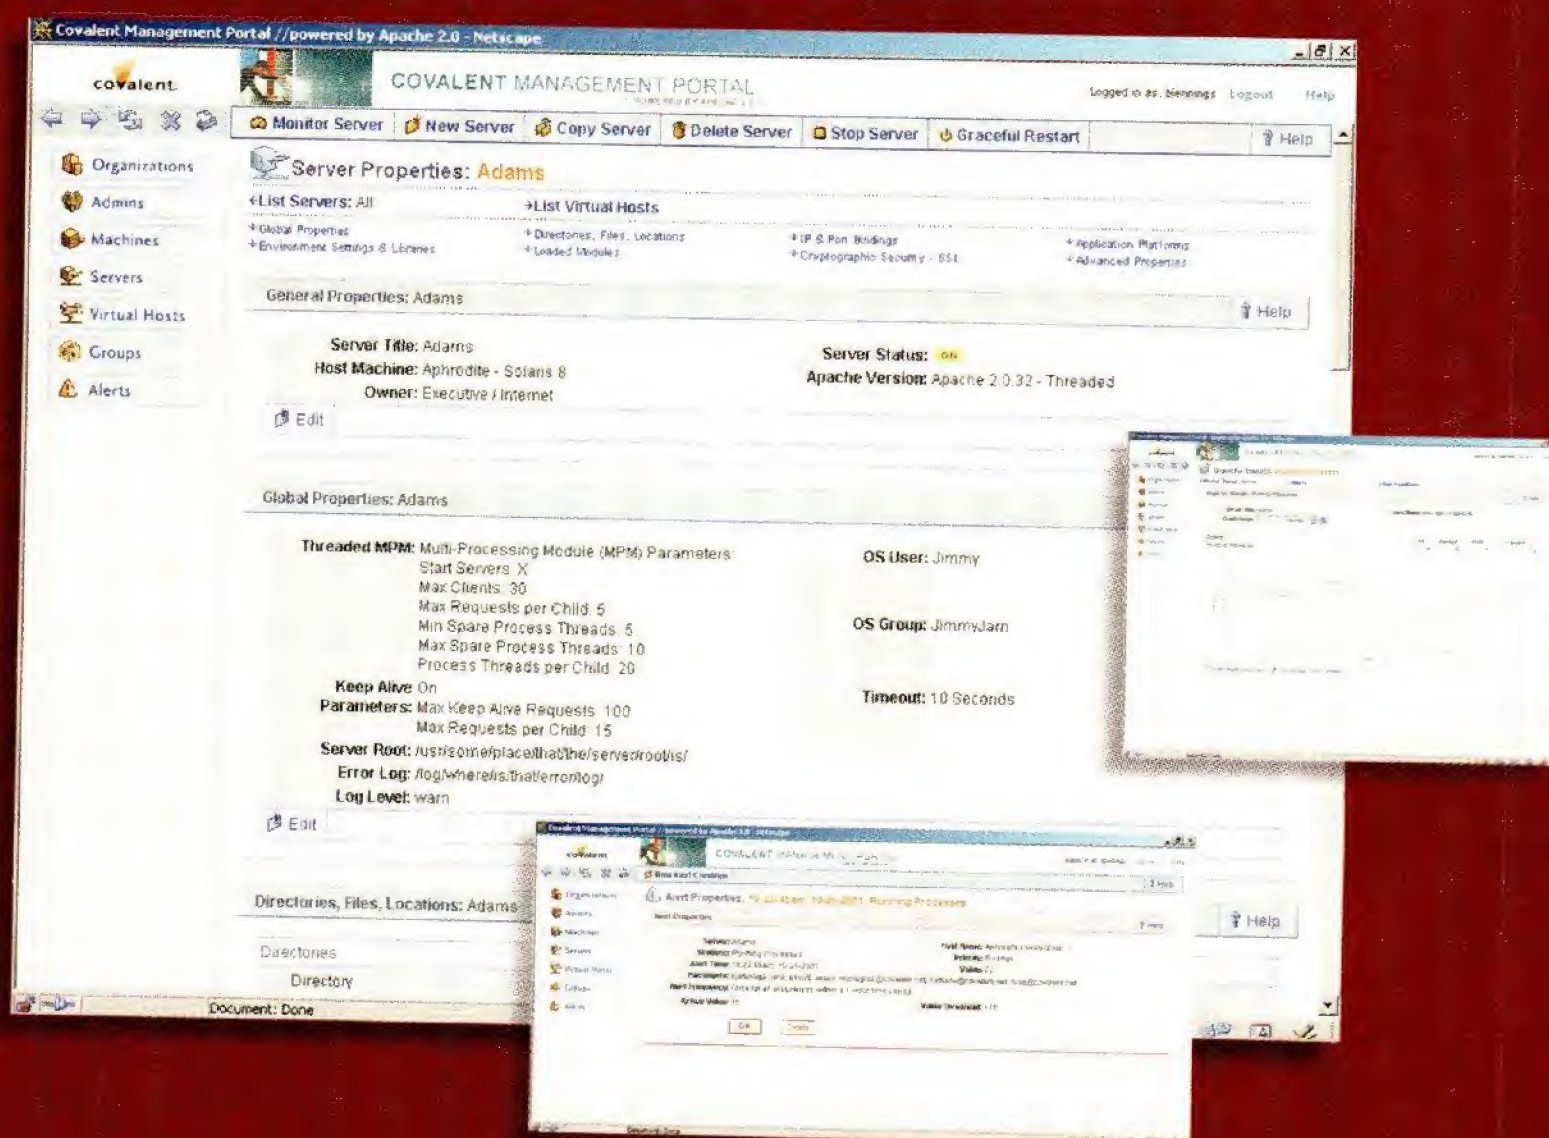
Finally, near the end of June, the SLOC found a rescuer in a company that had, in a sense, been by its side all along. It was NBC's joint venture with Microsoft: MSNBC.

MSNBC.com, the company that hosts and produces the popular news Web site of the same name, agreed to host and produce the Olympics sites. No financial details were made public, but industry buzz has it that MSNBC.com paid nothing for the rights, agreeing only to absorb the costs of the sites, reported to be around \$10 million. (Normally, an Olympics sponsorship goes for \$50 million and up.)

With fewer than five months before the planned November 2001 launch, and fewer than eight months before the opening of the Olympic Games, MSNBC.com had to build three sites—Olympics.com, NBCOlympics.com, and an Olympics version of its own MSNBC.com site. Together, the sites would contain over 5,700 pages of content and would have to support a bandwidth of 4Gb/sec.

"It was a short window in which to get everything done," said Mike Corrigan, director of technology for MSNBC.com. "We thought we could handle it, based on our experience publishing MSNBC.com on a daily basis." However, the Olympics sites posed challenges that the group had never experienced. For instance, MSNBC.com normally receives between 4 and 5 million unique visitors per day. The Olympic sites combined with MSNBC.com—all running on the same

FINALLY THE ONE THING APACHE HAS BEEN MISSING ALL ALONG



COVALENT ENTERPRISE READY SERVER POWERED BY APACHE 2.0

- Centralized browser-based management of hundreds of servers
- Enterprise grade Web server security
- Reduced total cost of ownership



FREE Find out what's new in Apache 2.0. Sign up for our free Web seminar at www.covalent.net/ers or call us at 800/444-1935 for more information on Enterprise Ready Server.

infrastructure—would, at their peak during the games, receive 8.8 million unique visitors per day. And because of the nature of the games, Olympics traffic often comes in surges, as news of triumphs, skullduggery, and controversy emerges. The group decided that the sites had to be able to handle a combined 300,000 concurrent visitors, compared with MSNBC.com's usual 35,000.

"At the beginning, this seemed nearly impossible," said Corrigan. "Having to design and build three sites...building an infrastructure for the largest sports event ever, and having to launch in five months and finish in eight seemed an incredibly tall order."

To top it off, the SLOC contract with MSNBC.com called for near-real-time results and online voting—two applications with more demanding timing and traffic characteristics than anything MSNBC.com had done before. For instance, the online voting application had to be able to handle 2.5 million votes in the space of a few minutes. Full results had to be online in less than five minutes, with critical events in less than one minute.

To ensure that the sites remained online and accessible in the midst of intense traffic, MSNBC.com planned to grow its physical infrastructure by about 40 percent, increasing the number of servers from around 30 to more than 50. The group added 100 contractors to its existing staff of 200. It also added a second data center to ensure more bandwidth and fault-tolerance in the form of redundant Web servers, application servers, and databases.

For load balancing between the two data centers, MSNBC.com bought two 3DNS systems from Seattle, WA-based F5 Networks. For local load balancing among servers within a data center, the group added four pairs of F5's BIG-IP controllers. The controllers are typically sold in pairs for redundancy.

glory or grief

Because Microsoft is one half of the MSNBC partnership, it isn't a coincidence that the sites would rely on Microsoft software. It marked the first time that Microsoft technology would be at the heart of Olympic Web systems.

At previous games with Web sites—in 1996, 1998, and 2000—the projects were completely IBM based. From the AIX operating system, RS/6000 SP Unix servers, and HTTP Server (a modified version of the open source Apache Web server), to the WebSphere Application Server Advanced, DB2 Universal Database, and Lotus Domino for content management, everything came from Big Blue. Even the load

balancing was done with IBM Network Dispatcher software.

To date, only Unix had proven that it could reliably scale to Olympic proportions. Microsoft had encountered limited success in going after the biggest sites on the Internet. Of the five most-trafficked sites on the Web (in descending order, according to www.netratings.com: AOL Time Warner, Yahoo, MSN, Microsoft, and Google), only the two controlled by Microsoft (MSN and Microsoft) use Microsoft technology. The other three use some variant of Unix.

Success at the Olympics would be a public-relations coup for Microsoft. Any significant glitch would strengthen the perception that Unix was king on the high end.

MSNBC.com planned to expand the infrastructure that supported the main MSNBC site to clusters of Compaq ProLiant 8500 servers. Each server had eight 700MHz Xeon processors and 4GB of RAM, and was running Internet Information Server, SQL Server 2000 Enterprise Edition, and Windows 2000 Datacenter Server.

the marathon begins

In late August, the group hired KeyLabs to start load testing the site. KeyLabs devoted three employees to the project full time. "It was a two-year project, and they had six months to get it done," said Mike Fahnert, president of KeyLabs.

KeyLabs started with a discovery process. A project manager and senior test leads came on site and talked to MSNBC.com to learn more about the traffic patterns and critical paths, as well as both current and projected requirements for scalability and load.

KeyLabs then began devising a test plan and test cases. Early on, KeyLabs decided to bring in the WebLoad testing tool from RadView Software. The WebLoad tool let KeyLabs simulate heavy traffic more efficiently, allowing it to drive more traffic from the lab to measure realistic loads on MSNBC's site. If the testing software had been less efficient, KeyLabs might have misinterpreted a

slowdown in the test machines as a slowdown on the MSNBC.com site.

learning from a setback

KeyLabs ran its first tests in early September. It was just in the preliminary stages when world events brought an unplanned stress test. The September 11 terrorist attacks pushed the MSNBC.com site well beyond any foreseen usage levels. Although servers didn't actually crash, they were only 43 percent available for about two hours following the attacks. During those two hours, the group quickly switched over to Akamai Technologies' EdgeSuite service, which caches entire pages on edge servers in Akamai's content delivery network. Normally, MSNBC.com serves only GIFs via the Akamai content delivery network (CDN). Switching the entire site over to the CDN let MSNBC.com remain online during the sudden traffic spike.

The switch to EdgeSuite took hours because MSNBC needed its servers configured and Akamai needed additional setup. Once these arrangements were in place, however, if a similar switchover ever became necessary again, it could be completed in minutes.

Akamai's EdgeSuite service was used in the Olympics disaster-recovery plan. "If there had been a major disaster, we could have turned on the Akamai EdgeSuite within a few minutes and hosted the Olympics site," said Corrigan. "This was our insurance policy." Although the system was tested, it was never needed for the Olympics.

The MSNBC.com group knew it couldn't miss its deadlines for the Olympics, despite the September 11 tragedy. It continued adding servers over the next few months. KeyLabs

SALT LAKE 2002
OFFICIAL SITE OF THE 2002 OLYMPIC WINTER GAMES

Sports Results/Schedules Athletes/Nations Spectator Info Shopping Games

Mostly Cloudy Tuesday, March 19, 2002
Current Salt Lake weather: 43°F(6°C)

Farewell, Salt Lake

- Reflections on a great Games
- Slide Show: The best performances of Salt Lake 2002
- Slide Show: The most inspirational athletes, stories

Medal Standings

	G	S	B	TOT
USA	12	16	7	35
Russia	10	13	11	34
Norway	11	7	6	24
Canada	8	3	8	17
Korea	2	4	10	16

- Complete list
- Medals by sport
- Medals by country
- Multiple medal winners

Shopping

OLYMPIC MEMORIES Salt Lake 2002 Commemorative Book
Seventeen days of breathtaking images and compelling stories

MSNBC.com's group of Olympic Web sites received 19.7 million unique visitors in 17 days.

developed and refined JavaScript-based test scripts, generating traffic from over 1,000 computers at its lab in Lindon, UT. To see how well the testing platform itself would scale as loads increased, the company verified script functionality, tuned and profiled the scripts, and examined the impact of scripts on the machines driving the loads.

In October, KeyLabs pushed the same scripts out to 240 systems in five data centers around the country. This system was capable of inducing the MSNBC.com site to transfer data at 2Gb/sec. Under the pressure of these tests, MSNBC became proficient at analyzing results and reconfiguring its back-end systems in hours rather than days. This was particularly necessary, as stress testing could only be performed for the approximately six off-peak hours each day.

In late November, KeyLabs pushed its test scripts out to a distributed testing network. It consisted of a million peer-to-peer nodes around the world and was managed by United Devices (UD).

The UD network is an established group of users who have loaded the UD agent on their systems. When a system is idle, the automated UD scheduler (jointly developed by KeyLabs and UD) can remotely initiate site testing. If a UD user starts using his or her PC, the scheduler immediately initiates a backup agent on another node.

the torch is fired up

Despite the monumental setbacks, the Olympics sites went live on November 19, 2001, the same day that Olympics officials went to Olympia, Greece to light the ceremonial torch. The sites were developed under budget and on time, said Corrigan, successes he attributes largely to constant and intensive testing.


Although a great deal of pre-games content was already online, MSNBC.com and the SLOC continued to add new articles and biographies to the sites from November 2001 until February 2002. They tested the live feed of the results data during November and December. The results system came online just one week before the start of the games.

They continued to upgrade the infrastructure with servers, load balancers, and caching

systems, and MSNBC.com also expanded into its second data center. At the end of December 2001, the group performed a final security sweep and lockdown of the site, searching for possible vulnerabilities to hackers, viruses, and other threats.

Finally, after eight months of pushing to get the sites, the infrastructure, the communication, and the back-up processes ready, the MSNBC.com team was ready for the Olympics. Friday morning, February 8, 2002 was an early day for the MSNBC team. By 6 a.m. most members were at work, nervously awaiting the first results from the K90 ski jump qualifying round, until the organizers rescheduled the event for February 10.

"[The cancellation] was a letdown for the team," says Corrigan. "Everyone wanted to



see the results system working in real time from Salt Lake, because if it didn't work it would be a significant blow to the site. Even though it had been tested many times, this was the real thing, with real time results, and no margin for error."

Everyone sat and waited and watched throughout the day. "There were teams in the data centers, teams in the Global Operations Center, teams down in Salt Lake, and teams in Redmond all watching intently to see what would happen," says Corrigan.

The traffic started to ramp up on Friday night with the opening ceremonies, but the staff was still nervous about the weekend, when the first real results would come in. On Saturday morning, most of the teams were in again by 6 a.m. to see how the results systems would work and watch for any other problems that might arise. At 7 a.m. the results started flowing. They were accurate, fast, and on the site in 15 seconds. The group breathed "a huge, collective sigh of relief," notes Corrigan.

the finish line

Before and during the Winter Games, fans accessed the Olympics sites in

record-breaking numbers, making the sports sites the most highly trafficked in history. Combined, they received 716 million page views in February and 19.7 million unique visitors over the 17 days of the Olympics. Usage for a single day peaked on February 21 at 8.8 million unique visitors and 74 million page views.

Metrics provider Keynote Systems noted periods of poor performance and unavailability on several other Olympics-related sites. For instance, the United States Olympic Committee (www.usoc.org) site had performance problems following Apolo Ohno's short-track speed-skating wipeout on February 16. Five days later, the site went down for nine hours as 16,000 emails protesting South Korean Kim Dong-Sung's short-track speed-skating disqualification rolled in.

The MSNBC sites were often battered at these times, too. Data compiled by webHancer indicates that on February 12, the traffic at Olympics.com increased by 37 percent, while NBCOlympics.com traffic went up 70 percent. At one point that day, 260,000 concurrent users were on MSNBC.com. But all the testing paid off. Not only did the servers stay up, but availability continued to average 99.8 percent.

The lowest availability at MSNBC.com was 99.5 percent on February 21, when Corrigan believes the site may have had 320,000 concurrent users. The Web team was not able to record concurrent users due to work it was doing on the voting application. But it did measure 1.3Gb/sec of actual egress (outgoing) data at one point in that day. Traffic was equivalent, in terms of page views and egress bandwidth, to the traffic the site experienced on September 11.

"This was a huge success for MSNBC, MSN, and Microsoft," says Corrigan. "In the end, the technology proved itself." The near real-time results were better than ever before, Corrigan adds—sometimes faster than television. The site was always available. In addition, the Society of Publication Designers selected the Olympics.com site for a Silver Medal design award out of 7,500 submissions.

As for plans for the next Olympics, Corrigan says that MSNBC.com is finally catching its breath. After an eight-month marathon, it's a well-deserved rest. ■

Michael (michael@hurwicz.com) is a freelance writer based in East Sound, Washington.

Biometric scanners may soon be a real alternative to weak, password-based security. But are fingerprint readers and iris scanners sound investments?

critical decisions

Jerri L. Ledford

the rolls royce of security

Are biometrics worth the expense?

Just five years ago, most application developers regarded fingerprint scanners and voice print analyzers as the stuff of science fiction. The devices made for good security, but the cost of such high-tech gadgetry made it beyond the reach of all but the most dedicated mad scientists.

For the last few years, however, biometric security devices have been a growth area in the technology industry. Pharmaceutical companies, hospitals, financial institutions, and corporations are all beginning to move from using simple password security to something far more secure—a unique identifier, one that's very hard to fake. As a result, there's probably a biometric device coming to a network or application near you.

As the applications for biometrics have broadened, the price of biometric scanning hardware has decreased, and is expected to keep dropping. The International Biometric Group (IBG), a research firm in New York City, expects revenues from biometrics products to grow from \$399 million in the year 2000 to \$1.9 billion by 2005. So is biometric security right for your application? There are several factors to consider.

outdated security

The Internet and wide-area networks have lengthened the reach of many IT organizations, but they've also made networks more vulnerable. Frustration with current security technologies has been growing exponentially. Each time a secure system is compromised it

can cost a company hundreds of thousands of dollars in damages and downtime, making inadequate security a serious risk for any organization.

Hackers have compromised even some of the more elaborate hardware-based authentication systems. Take, for example, the case of 36-year-old Serge Humpich, the French hacker dubbed "the Count of Monte Crypto." In February 2000, Humpich used a \$250 piece of equipment to crack the encryption key embedded in the Smart Chip-enabled debit cards used by some 22 million French consumers.

Humpich claims it took him four years to defeat the security system used in the French bankcards, but then, that system used 640-bit encryption. Less sophisticated designs, like the one used for 802.11b wireless networking, have been cracked more easily. If an intruder wants into a network badly enough, given the proper equipment, a little know-how, and enough time, he or she can access any system secured solely by encryption.

Traditional security methods are vulnerable to still other risks. Often, it's simple laziness that causes employees to compromise password-based security. "It is very common to see people posting or sharing

their passwords," says Santosh Chitakki, vice president of marketing for Vienna, VA-based BioNetrix, a biometrics software provider. "Biometrics are a very secure way of enabling users to access applications," Chitakki says, noting that biometrics replaces the password with a key that can't be shared.

which device?

Biometric security works on a simple principle: no two people are built exactly the same. Whether the actual process involves scanning a body part or imprinting voice patterns, biometric software maps physical characteristics, called "markers," that exist in a combination that's unique to each individual. The three methods used most often are fingerprint scans, iris scans, and voice print recognition.

Voice printing is considered the least secure and least user-friendly of the three. Keith O'Leary, the United States marketing director for Keyware, a European biometric-enabled applications vendor, says there are some problems with voice printing. For example, outside factors like throat inflammation can cause false readings.

Fingerprint scanning is the form of biometric security in widest use today. When

"The most expensive part of deploying biometrics is making sure the user population is trained, and making the enrollment process as easy as possible."

a user places his or her finger or thumb on a reading device, the device scans the print and then matches it to a mathematical representation of the known fingerprint for that user. The mathematical representation is built from over 100 identifiable markers found on each fingerprint. At least 70 of these markers must match for authentication to succeed.

Iris scans, on the other hand, rely on 266 points of recognition within the iris of the eye. If 150 of those points match, then identification is conclusive. "Iris recognition is changing the biometric landscape," says O'Leary. "With each point that matches over 150, the ability to conclusively verify a [person's identity] is more accurate."

O'Leary explains that fingerprinting is enough in most cases, but it all depends on the company that is deploying the technology. "What is your security need?" he asks. "What is your methodology, and what will be the level of security that you want?" In some cases, he says, a single biometric technology might be used at a lower level of security, with additional biometrics added as the user accesses more-critical applications.


name your price

Until recently, one of the biggest barriers to biometric security has been the cost of implementing a system. In the past, the reading devices alone could cost thousands of dollars, and that didn't include the software or technical expertise necessary to use them. The total price tag per application was well over \$150,000. While many organizations wanted a more secure method of protection, the return on investment was hard to justify.

Today, the tables are turning, and device costs are dropping dramatically. For example, many computers currently come equipped with a microphone, and the Mac OS even ships with voice recognition technologies built in. To upgrade to a noise-canceling microphone like those made by Santa Cruz, CA-based Plantronics typically costs no more than \$50.

Similarly, some fingerprint scanners cost as little as \$50 per device, and industry analysts predict that as biometrics gain in popularity, that figure will fall even further. Both Dell and Compaq already offer add-on fingerprint scanning pads as an option on desktop computers, and Acer offers laptops with those devices as standard equipment. IBG predicts that by 2005 every computer will ship with biometric devices.

By comparison, BioNetrix's Chitakki estimates that it costs about \$15 per call to help a user who has lost or forgotten a password. Biometrics can present an extremely cost-effective alternative to traditional security. "[Administrators] end



up gaining convenience and a reduction in administrative costs," says Chitakki.

Overcoming users' initial resistance to biometric systems can be a problem. But once biometric scanners are incorporated into devices that we already use, such as the mouse, users will warm to the technology more readily. What's more, middleware software providers like BioNetrix and Keyware, which reduce the need for in-house technical knowledge and biometric systems, are becoming attainable for any size organization.

standards

"The most expensive part of deploying biometrics is making sure the user population is trained, and making the enrollment process as easy as possible," says BioNetrix's Chitakki. Every time biometric security is applied to an application, that application's existing users must re-enroll with a new fingerprint, voice, or iris scan. This can become tedious if several applications require authentication, and can lead to inconsistencies. However, once single sign-on services like Microsoft Passport or the Liberty Alliance project become more commonplace, they will likely help to ease such difficulties.

Other than cost, the most important consideration for application integrators who are considering a biometric security option will be the efficacy and life cycle of the hardware involved. The ease with which biometrics can now be added to a network shouldn't be confused with the quality of the biometric device or the standards that might affect the devices and software in the future.

"This is a changing technology," says Chitakki. "You need to simplify your security plan, plan ahead, and buy middleware software that reduces the number of times that you have to inconvenience users. You also need to be sure that the biometric

systems you implement will support multiple applications and devices."

Choosing a product that supports a wide range of applications is important, agrees Keyware's O'Leary. One shortcoming of the biometric security industry is that it still lacks standards. Perceiving an open playing field, numerous companies have begun to offer biometric security solutions with

little concern for interoperability. Among these are the BioAPI Consortium, the National Institute of Standards and Technology (NIST), and the Security Industry Association (SIA).

Although it isn't uncommon to have this disparity in standards in a young industry, O'Leary says that "Lack of standards from a risk assessment perspective should always be a concern. You should look to the people who are forerunners and who are participating in the standards bodies."

what's to come

"It's a new technology, and we're learning," says Chitakki. Yet he feels that his company and others in the industry are rapidly bringing biometric security to the point at which it will become attractive to a wide range of organizations. "Prices are dropping, form factors are changing and [the devices are] becoming unobtrusive, and the improvement in quality is incredible," he says.

Even today, adding biometrics may be worth the expense if it helps to improve security and reduce workflow costs. But if both of those conditions can't be met, then it might be best to wait a little while longer. Ongoing technical improvements, coupled with the increasing need to safeguard information and applications, are likely to transform biometrics from the Rolls Royce of security to the Honda of everyday protection. ■

Jerri is a freelance technology writer. You can contact her at jerriledford@cs.com.

Desktop application programmers have been using Model View Controller design for years to improve their development processes. Applying the same tools to Web development is easier than you think.

critical decisions

Al Williams

design patterns for web programming Do you need MVC?

Between consulting and training, I get to glimpse the inner workings of many different companies and their development projects. Recently, I've noticed that as the demands placed on Web applications increase, some applications aren't up to what's now required of them.

The problem is their simple architectures. Although unsophisticated designs are sufficient for basic, single-purpose Web sites, these same programs often fail when they're extended into new territories. For example, a simple set of Web pages and CGI scripts isn't always a good fit for, say, internationalization. Bundling data into a form that's suitable for VoiceXML or display on a mobile phone isn't trivial either.

Traditional development teams wrestled with similar issues when they made the transition from text-based applications to graphical programs under Windows, the Mac OS, and the X Window System for Unix. The solution that many developers turned to was the Model View Controller (MVC) design pattern.

Many programmers will be familiar with MVC from popular C++ application frameworks like the Microsoft Foundation Classes. I first encountered MVC when I was learning SmallTalk, and I had a lot of trouble with it initially. However, over time I've decided that it really does make sense for interactive applications.

All the same, MVC hasn't yet been widely adopted in Web development. MVC tools for the Web exist; a few frameworks are available, including Maverick (mav.sourceforge.net) and Struts (jakarta.apache.org/struts). But like

most design patterns, you don't really need a special tool to use MVC (although, a pre-built tool would probably be easier to use than rolling your own solution). So why hasn't MVC overtaken the Internet the way it has classical application development?

Some developers I talk to don't have a clear idea of what MVC really is. Others think MVC is too much trouble for a Web application. Of course, for a simple page, practically any scripting is overkill. But if you have anything that resembles an application—that is, you interact with the user, accepting input and generating results—MVC deserves your consideration.

what is mvc?

In a traditional application, a single piece of code handles everything. With MVC, you break down your program into three cooperating parts: the Model, the View, and the Controller. The View is the part the user sees. It formats data into an onscreen representation. However, it doesn't actually contain the data. The data resides in the Model. Finally, the Controller portion accepts user commands and modifies the Model.

Naturally, the Model has to inform the View that it should update the representation of the data when that data changes.

The classic example of this strategy is a spreadsheet program that runs on a personal computer. Using MVC architecture, the Model stores the formulae and other data. When you issue a command to save to a file (or load from a file), the Model handles this action. It also handles the specific logic, like recalculating the entire sheet. The View draws the familiar grid that shows a part of the data (depending on the scroll bar's position). The Controller deals with any process in which the user changes something.

This approach has several advantages. The most obvious is that this decoupling allows for easier unit testing. It also isolates changes to a large degree. Changing the way data is displayed is not likely to affect the document object functions. However, an even greater benefit is the ability to mix and match different parts. For example, you might want to add a pie chart View of the data. Not only is this relatively easy, but you can reuse the pie chart View with other applications that use the same style of document object.

mvc and the web

In the JSP and Java world, MVC development is often referred to as Model 2. Java developers sometimes use MVC by creating a servlet to

The bottom line is that MVC applications are more flexible and more able to withstand change than most ad hoc solutions.

act as the Controller. All requests go to the Controller servlet, which then decides which JSP to invoke to process and display the results. The JSP acts as the View, and *link* or *submit* buttons within the JSP play the role of Controller. If you aren't comfortable with writing servlets directly, I've seen people write Controllers using straight JSP. (After all, a JSP compiles into a servlet).

Java isn't the only way for Web developers to take advantage of MVC. ASP.Net has some MVC-specific features built into it, as does Apple's WebObjects, but you can use MVC with practically any scripting or programming language. You could even use Perl. For example, the Perl scripts at eXtropia (www.extropia.com) use a custom MVC framework. In fact, you can adapt MVC to any of the popular scripting languages. You simply have to partition your code so that one part handles the underlying data, a second part handles the presentation, and a third part handles the changes and control.

You might also consider using XML to implement your Model. Because the Model part of MVC is supposed to hold the data, it makes sense that modern MVC architectures would work well with XML. Decoupling the View and Controller logic lets the document code focus on the XML parsing. The Views can just assume that they have the data they need without dealing with the intricacies of XML. Changes in the XML structure only affect the Model code. Of course, with the move to supply XML-based Web services, your document might even be distributed between multiple servers with little difficulty.

XML is a great choice for many applications. Of course, one of the nice things about MVC is that with proper design you could use XML, a database, or even an ordinary file for the data source. If the interface to the rest of the program remains immutable, no one will care where the data originates (except, of course, for the person responsible for the Model).

an asp example

Naturally, if your language of choice includes support for MVC, you can expect your program design to be more elegant than if you

had to roll your own support. However, MVC doesn't require any particular support from your scripting language. It's more of a philosophy than a specific technique.

To illustrate, I've designed a simple MVC framework using XML that's written in classic ASP. The source code for this framework is online at www.newarchitectmag.com. Naturally, there are many ways you might use ASP to implement an MVC strategy, but this simple example is easy to follow and will give you a concrete example of MVC in action.

Even though ASP isn't designed with MVC in mind, the `Server.Execute` and `Server.Transfer` methods make it simple for one script to call or jump to another. A request to the controller might look like this:

```
www.some-web-site.com/control.asp?
cmd=viewtable&doc=intlist.asp
```

The `control.asp` file serves as Controller, and as an entry point to the entire application. By providing different commands and documents, this simple script reads the document data and passes control to the appropriate View. In addition, the Controller provides defaults for the parameters (`cmd` and `doc`) if you don't provide any values.

ASP isn't well-suited for passing data between separate modules, so the document object reads the XML data (using the Microsoft XML parser) and stores the resulting information in the current Session object. As a side benefit, this action also caches the data, so the document object only reparses the XML when necessary. If you have a lot of data, you might have to flip everything

around and have the

Controller call the View. The View could then call the document and request a particular subset of data. That isn't a problem. MVC requires that each part of the system be responsible for managing its own domain (the data,

the control, or the display), but that doesn't mean that the responsible piece can't delegate to another module.

The View, of course, simply formats the data it finds in the Session object. It doesn't know, or care, where the data was before. I'm not seriously suggesting that you'd use this simple-minded example as a skeleton for serious applications. I merely want to illustrate that you can implement MVC in a mere handful of ASP lines. Thus, you should be able to adopt it for practically any project.

because it's there?

As Web applications become more complex, the advantages of more complicated development techniques like MVC become increasingly evident. The benefits are numerous:

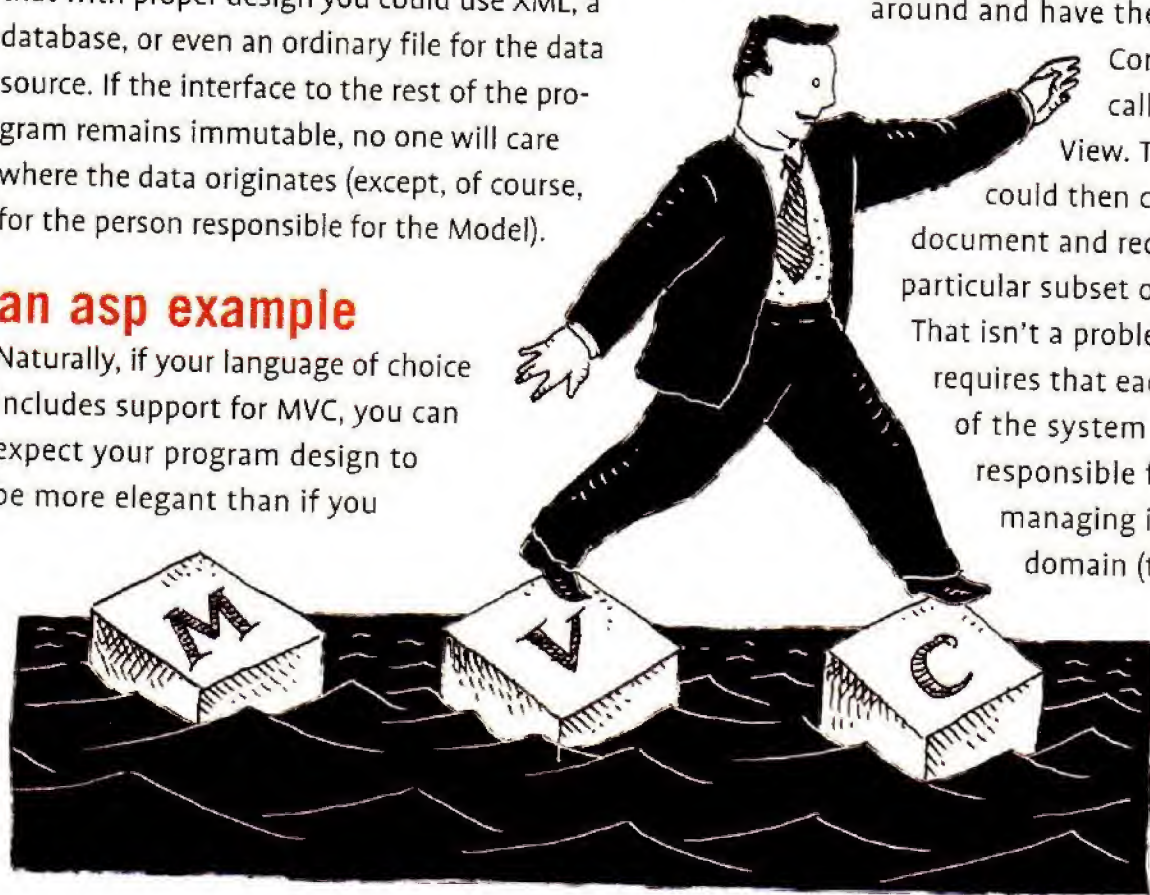
- MVC simplifies the creation of multiple user interfaces with the same data.
- It increases the opportunity for code reuse.
- It lets developers write and debug modules independently.
- It decouples code, allowing for easier modifications and division of labor between client and server.

The bottom line is that MVC applications are more flexible and more able to withstand change than most ad hoc solutions. Of course, simply because you can do something doesn't mean that you should. But for any application beyond a simple page, you often receive compelling benefits using MVC.

MVC software may take a little extra development effort initially, but that quickly pays for itself in decreased debugging time and increased flexibility. As a bonus, MVC increases the possibility of code reuse, which can further reduce total costs over the life of several projects.

As I've demonstrated, MVC is a methodology you can apply with nearly any development tool. Of course, if you have the luxury of choosing new development tools, you might want to consider selecting one that provides direct support for MVC and makes your life even easier. But in the end, regardless of what you use, MVC can help you design architectures that will withstand the test of time. ■

Al is the author of many popular programming books including Java 2 Network Protocols Black Book (Coriolis). You can find him on the Web at www.al-williams.com.



STRATEGY

Industry leaders talk about innovation and the future of the Internet.

DEVELOPMENT

XML-based standards bring the Semantic Web one step closer to reality. p.30

INFRASTRUCTURE

Eliminating server management headaches with autonomic computing. p.34

**COVER
STORY**

by *New Architect Editors*

by design

wisdom from the industry

WE ASK 16 OF YOUR PEERS ABOUT THE TECHNOLOGIES AND INNOVATIONS THAT ARE CHANGING THEIR JOBS.

**Charles F. Goldfarb**

Inventor of SGML


What led to the creation of SGML?

Before SGML there was IBM's GML—not standardized. It was created in 1969 to solve what today we would call an application integration problem. I was trying to get three separate programs to work together to form a law office information system—an editor, a formatter, and an information retrieval program. I hit on generalized markup as the means to let them share data. With Ed Mosher and Ray Lorie, I developed that idea into GML.

I continued research on document structures and created additional concepts, notably the validating parser in 1974. At that point, SGML was born, although it took another decade before it was fully developed and standardized.

What would you have done differently?

I'd have documented it better, especially the use of notations for strong datatyping and the relationship between documents and data.



How will we reach the goal of truly separating presentation from data?

We've seen it happen in most applications of XML because XML is used chiefly for data interchange and messaging. Those areas emphasize the processing of the abstract data, so rendition is normally handled with a separate style sheet. But browsers and traditional rendered Web pages are, of course, another story. I think we'll start to see an improvement as we ship more XML directly to the desktop, rather than converting it to HTML on the server. Now that the major browsers can support XML, that trend will increase.

Content management systems are also in their early phases and known for their cantankerousness. What will the future hold for such systems?

As enterprises increasingly recognize the importance of content management, content management systems will be layered on DBMSs and application servers. They're likely to spur the growth of native XML DBMSs and XML layers on relational DBMSs.

Robert Hopkins
Founder, Weblations

What's wrong with automatic translation tools?

Machine translation (MT) tools simply don't work very well. To sum up recent progress somewhat cynically, I would say that today you can get a bad machine translation in only a fraction of the time it took 15 years ago.

MT will someday be improved through new algorithms based on the latest research into how the brain processes language. Meanwhile, computing power is increasing. Researchers are building huge databases of translated material to feed and test the new algorithms.

How is Unicode affecting the World Wide Web? Is it making multilingual Web sites easier to build and maintain?

My company specializes in Web site localization. The Unicode revolution has brought huge benefits to our tools and platforms, but strangely, it's non-existent on the Web itself.

When you browse a page on the Web in English, Chinese, or whatever, the chance that it's delivered in Unicode or UTF-8 (the Internet's compact version of Unicode) is less than one in a million, literally.

What's the biggest obstacle to internationalization, and how will it be overcome?

The biggest obstacle to internationalization is the sheer distance, in every sense of the word, between the content owner and the content user. The people who benefit from I18N mostly are not Americans, while those who pay for it mostly are.

Most of the biggest knowledge bases are available in English only. Why? I guess it's because the Americans with the big budgets don't understand in their bones how absolutely imperative it is to translate that material to get a return on it outside of the U.S.A.

Richard Luna
CTO, PasswordHeaven.com

What's the number one security problem with the Internet, and how will we fix it?

A lot of areas are susceptible to attack, but one area touches most of us: email. For all of its benefits, email is clear text. The vulnerability is compounded by TCP/IP.

The answer is encryption. Some email is encrypted today, but unfortunately most devices can't access the encrypted email, and therefore, the standard is a non-standard.

Microsoft has said it is rededicating itself to security. Will it succeed?

No. Microsoft's primary focus is increasing its dominance in the market. Strengthening security doesn't help in that area.

One example is what it did with the XP TCP/IP stack. In previous versions of Windows, the core IP address couldn't be spoofed. That meant that if an attacker was using a Windows-based PC, the attack could be traced back. With XP, Microsoft changed the stack so that the originating IP address can be spoofed, making XP a good hacker PC. If Microsoft were serious about security, it would have left the TCP/IP stack alone.

Is biometrics something to watch?

For biometrics to be beneficial, it has to work across all operating systems, not just one or two. Market forces will prevent that from occurring.

Would Microsoft ever let a mainframe company or Sun dictate its authentication? All of these vendors know that owning the authenti-

U.S. forms the Advanced Research Projects Agency (ARPA).

1958

Packet-switching network presented to ARPA.

1961

Leonard Kleinrock of MIT publishes the first paper on packet-switching.

1968

1969

UCLA, UCSB, Stanford Research Institute, and University of Utah set up first four nodes on ARPAnet.

Ray Tomlinson of BBN invents the email program.

1971

1972

First computer-to-computer chat takes place at UCLA.

cation has the effect of owning the enterprise. All of them, from Novell to Microsoft, Sun, and IBM, are trying to own the enterprise. Because there are so many vendors competing for authentication ownership, no one group will win, and therefore biometrics will be relegated to small or platform-restricted applications.

Also, the authenticating system has to have a pattern to match with. Without the original pattern, the authentication mechanism fails. With biometrics, the authenticating pattern is your fingerprint, palm print, retina, etc. I don't believe end users will give their consent to store their authentication information.

What do you see on the horizon for encryption?

Encryption is the future. Whether we're talking about VPN communications or email, encryption is the only mechanism that makes Web use truly secure.

As more systems are interconnected, encryption will be required to keep the communication lines from being disrupted. That said, before encryption gets fully deployed, I think—unfortunately—we will see several major disasters. Just as the September 11 attack on the World Trade Center reminded the IT world not to concentrate technology in one location, managers will need to experience an attack [before they will] provide financial resources for encryption.




Photo by Webb Chappell.

Henry Lieberman

Agents Research Lead, MIT Media Lab

What are the main problems with today's technologies that agents can help solve?

Complexity. We think of a computer today as like a box of tools. Applications and menu operations are like hammers and screwdrivers. Each is specialized to do a particular task. There's nothing wrong with tools, but having a different tool for each task means that if you want to do too many things, you get too many tools, and it takes too many steps to do anything.

The alternative is to cast the computer in the role of an assistant or agent, like a travel agent, a secretary, or a real estate agent. Computer agents don't have to be as smart as a person, but they do have to be proactive. They do have to learn from interaction with the user, and they do have to be sensitive to context. That's what we're working on.

Why haven't agents taken on a more prominent role, especially now that Internet technologies are making it easier for people to connect and collaborate?

Industry is adopting agents, but in small ways, and slowly. Amazon recommends books it thinks you might like. American Express uses an AI program to approve credit purchases. There are many such small stories.

But industry sticks to the idea of tool-based interfaces for the most part, and treats agents as an add-on. It should redesign interfaces from scratch to take advantage of agents.

What sorts of tools will end users need in order to take advantage of agents in everyday computing?

One of the things I think is most important is to give ordinary end users the ability to teach the system new behavior. Every user or business has unique needs and desires, and no company can build custom applications for everyone.

We can do this through a technology called *programming by example*, where the user demonstrates concrete examples of what they would like the computer to do, and a software agent records the examples and makes a generalized program out of it. You can think of this as macros on steroids.

What does the future look like for human-computer interaction, and what do technologists need to do to get there?

Certainly there will be many hardware developments that will transform the interface, such as new input devices and display technologies. Speech interfaces will be big. In software, I think we have to break out of the desktop/application/tool metaphor, and agents will be a big part of this transformation. There will be a consumer revolt against technologies that are too complicated, unreliable, or constraining, in favor of technologies that are simple and elegant, truly helpful, mindful of context, and flexible enough to meet a wide variety of needs.

William (Tilt) Thompkins, Jr.

Persistent Virtual Spaces Research Lead, National Center for Supercomputing Applications

What is a Persistent Virtual Space?

Virtual spaces use tools such as chat, application sharing, and confer-

- Robert Metcalfe outlines the idea for Ethernet.
- Vinton Cerf and Robert Kahn present basic Internet ideas at the University of Sussex.
- File Transfer Protocol (FTP) refined by A. McKenzie.

1973

1974

- Cerf and Kahn publish "A Protocol for Packet Network Interconnection" detailing TCP.
- SGML invented by Charles F. Goldfarb.

Unix-to-Unix Copy (UUCP)
developed at AT&T Bell Labs.

1976

1978

TCP split into TCP and IP.

1978

Usenet established between
Duke and UNC by Tom Truscott,
Jim Ellis, and Steve Bellovin.

encing to assist real-time or synchronous collaboration. The idea behind Persistent Virtual Spaces is to add recording, indexing, and archiving tools to support asynchronous collaboration. This dimension provides an interaction that isn't possible with face-to-face meetings.

What security and privacy concerns are there with storing and using data about people's interactions?

Privacy is a key concern anytime information is recorded about people, whether it's personal data, email, or conferencing interactions. I don't think Internet users are aware of how insecure transmission is, but there is growing awareness of how persistent and open to discovery email conversations can be. Strong encryption of stored data as well as transmissions is going to be required to maintain privacy.

What technologies will it take to get us to a point at which Persistent Virtual Spaces are a reality for companies and even households?

We find that voice interactions play a key role in driving group interactions, and quality voice-over IP seems to require around 1Mb of bandwidth. Until this kind of bandwidth—fixed and mobile—is available to participants, I don't think virtual space usage will take off. Once the enabling voice threshold is crossed, several support items like network attached storage and encryption become critical. As users begin to collect numerous interactions, I believe indexing, cataloging, and retrieving those interactions becomes the most difficult issue—[one] which drove us to focus on object-oriented metadata cataloging.

What other innovations do you see coming from the academic sector over the next five years?

Two areas where the academic sector has clear interest and momentum: seamless connectivity across fixed and mobile interactions, and deployment of microsensors in office and external environments.

Tom Frobase

Manager of Supervisory Control and Data Acquisition & Telecommunications, TEPPCO Partners

How will engineers sift through increasing amounts of data?

As the data volume increases, so does the management of that data. Broker-based tools like IBM's MQ series, Microsoft's MSMQ, and

Sonic's SonicMQ will replace some of the massive stores of data used today. Random access to pseudo real-time data will increase the flexibility of modeling real-time processes, while diminishing the need to deal with irrelevant data.

Will new installations take advantage of networking differently than those built before the Internet was a major force?

The energy industry has billions of dollars worth of equipment currently installed, with many years of useful life remaining. The common approach is to create embedded adapters, providing interfaces from legacy hardware and serial protocols to high-level XML structures using a Java Messaging Service interface. This lets the legacy hardware function in the future environment, while providing a methodology for equipment makers to follow as they supply replacement equipment [that has] direct XML TCP/IP interfaces. We have recognized that we cannot wait for our suppliers to offer network-ready equipment. The energy industry is providing an infrastructure that supports networking, and with that, supplies a common framework for the instrumentation vendors.



Regan Moore

Associate Director of Data-Intensive Computing, San Diego Supercomputer Center

How are data grids being used today?

Data grids provide a way to assemble and manage collections of data that reside at multiple sites. The collections are used by a community to make it easier to share results, gain access to experimental data, or facilitate collaborations. Effectively, data grids provide a way to share data without having to worry about differences in types of storage systems, security environments, or access mechanisms.

Data grids provide name transparency (find a digital entity without knowing its file name), location transparency (find a digital entity without knowing where it's stored), and access transparency (retrieve a digital entity using the same API from archives, file systems, databases, and FTP sites). The technology we use to accomplish this is the SDSC Storage Resource Broker and Metadata Catalog.

Name server developed at University of Wisconsin.

1983

Symbolics.com becomes first registered domain name.

1984

Domain Name System (DNS) introduced.

1985

1986

- Network News Transfer Protocol (NNTP) proposed by Brian Kantor and Phil Lapsley.
- Mail Exchanger (MX) records developed by Craig Partridge.

Perl released by Larry Wall.

1987

ATTENTION RESELLERS

web hosting

It's 3 o'clock in the morning and you're still waiting for an answer to the 4 emails you sent Tech Support. No one has returned the cries of desperation you left on the answering service. The 24 hour support they supposedly offer is going to make you look bad tomorrow morning when your client does not understand why you can't deliver the products and services you've promised. Where do you turn when you're alone in the dark and your head is sore from banging it against the wall?

At AIT, we know how you feel. We used to host our business with other companies. In fact, we used to be resellers. Our business depended upon the quality of service provided by our hosting partners. We've built AIT from experiences like yours and understand your needs. Five years ago we were building and operating websites at 3 am out of a garage. We've grown into a 90,000 sq. ft. facility and data center and are 200 employees strong. We haven't forgotten how frustrating it is and run a business when no one seems to understand how crucial it is that you get some answers to your questions.

If you're looking for a hosting company that understands what you need and cares about your success, then pick up the phone. We'll answer it. We'll help you get your sites online and working, even at 3 o'clock in the morning.

We don't take your success for granted. We give you the tools to succeed. We keep our prices affordable so you can make a profit. Your success really depends on only 3 things: cost, products and services you offer, and support for those products and services. **We guarantee that no other hosting company offers more features at better prices than AIT.** What are you waiting for?

CEB

Clarence E. Briggs III
CEO, Founder



ADVANCED INTERNET TECHNOLOGIES
& COMMUNICATIONS, INC.



Read how I did it at <http://aitcom.net/story>

We give you the tools to succeed. Use them yourself and resell to all of your clients for a profit.

Website Design Tool
FREE Shopping Carts
FREE Secure Directories
FREE PAL Membership
eTools Office Suite (Auction, Chat,
Web Based Email, BBS, Calendar, Banner Exchange)
Unlimited POP Emails
Servers Enabled for Wireless

AIT offers a wide range of hosting solutions - Personal/Business/Reseller/Corporate

Your Hosting Solution	Small Business Plan (Server 1)	Standard Business Plan (Server 2)	Advanced Business Plan (Server 3)	Corporate Business Plan (Server 4)
Disk Space (can be customized)	225 MB	450 MB	850 MB	30 GB
Max # of domains you can host	N/A	30	85	300
Base Monthly Server Cost*	\$15.95	\$49.95	\$129.95	\$319.95
Max. per domain cost @ \$3.99/domain/mo	N/A	\$119.70	\$339.15	\$1,197.00
Max. possible cost to you/ month	\$15.95	\$169.65	\$469.10	\$1,516.95
Your monthly gross profit @ \$21.95 domain	N/A	\$658.50	\$1,865.75	\$6,585.00
Your monthly net profit reselling hosting	N/A	\$488.85	\$1,396.65	\$5,068.05
Additional Profit Reselling AIT Extras**	N/A	\$1,500.00	\$4,250.00	\$15,000.00
Total Monthly Profit	N/A	\$1,988.85	\$5,646.65	\$20,068.05

*Requires 24 month contract

**Assumes Reseller sells the many FREE AIT offerings at \$50/month per customer (Standard Accounts)

1-800-878-4084

AITcom.net
sales@aitcom.net

As we put more data online, how will we store and manage it?

Archival capacity is growing exponentially. The archival storage system currently has a capacity of 900TB. With next-generation tape technology, the capacity will grow to 9,000TB or 9PB. We can expect to see a doubling of storage capacity every one and a half to two years.

The ability to discover relevant data will require the ability to manage semantic relationships between collections. Knowledge management is the current area of intense research.

And how will we deal with the threat that an attack, a glitch, or a virus might cause us to lose entire libraries at once?

Data grids support replication of digital entities across multiple sites. Thus, if data is lost at one site, it can be retrieved from another. The critical element in any data management system is the catalog that holds the information stating where the digital entity is stored. A well-engineered system provides multiple ways to safeguard this information.

Monika Henzinger

Director of Research, Google

It has been said that innovation is a process, not an accident. What does Google do to keep innovative ideas flowing?

Everybody is encouraged to come up with new ideas, and we have a Web page where everyone can post these ideas. There's a meeting every two weeks to discuss them, and everyone is invited.

The Google News Search came out of these meetings. It was one of these topics that we had discussed for a while, like "Wouldn't it be great if we could do something with the latest news?" And then we came up with the idea to crawl newspapers and their latest stories.

Which Internet technology has most impacted business?

Email. It has connected everyone.

Which technologies could you have done without?

I interpret this more as, "What technologies would I be glad to have never seen?" Pop-up ads. They're completely against Google's philosophy. Technology should help the user, not annoy them. I don't go to Web sites anymore that have pop-ups. As Google shows, you can be successful without them.

How does your experience with WAP contradict (or match up with) what critics say about it?

People criticize WAP because of the small screens and all the typing you have to do. But i-Mode works very well in Japan, with similar limitations. At Google, we have built *nav detection*, so when we send a page with WAP, we send only one navigation link to reduce screen clutter. You can click on the link to get to the rest of the navigation. For the problem of having to press number keys many times to get the proper letter, *number search* lets you type in a search query just by pressing each number once. From the numbers, Google will figure out what word you're trying to type.



Donald Norman

Co-Founder, Nielsen Norman Group

You point out that many devices and technologies are poorly designed with respect to human factors. What's the main reason for this, and what can we do about it?

The factors that make for success in the marketplace are still being debated. Ease of use, traditionally, was never high on the list of concerns, but as electronics started creeping into our lives, these issues grew.

All of the best companies hire excellent designers with human-centered design skills. But as fast as things get better in one industry, they deteriorate in others. DVDs have unusable menus, and when we complain, we're told that this is artistic freedom. Well, I want artistic creativity in the movie, not in the things I have to do to play the movie.

The weird thing is that companies lose money for lack of attention to these issues. Service calls are expensive, yet every company now has to man very large, very expensive call centers to deal with perplexed customers. Moreover, some customers have just tuned out. Why buy a new computer and go through the hassle of upgrading and change when the old one works just fine?

In your opinion, which devices over the past two decades have truly been innovative?

The Macintosh computer. The Web browser. Portable CD/DVD players. MP3 players that fit in small pockets. The cell phone and the way it has come to dominate our lives.

1988

- Computer Emergency Response Team (CERT) formed by DARPA.
- Internet Relay Chat (IRC) developed by Jarkko Oikarinen.

1990

- The World (world.std.com) becomes the first commercial provider of Internet dial-up access.
- HTTP 1.0 specification published.
- The first HTML document published by Tim Berners-Lee.
- Archie file-indexing system developed.

1991

- World Wide Web released by CERN and Tim Berners-Lee.
- Gopher released by Paul Lindner and Mark P. McCahill from the University of Minnesota.
- Linus Torvalds announces release of the source code to a Unix-like kernel (to become Linux).

1992

Veronica, a Gopherspace search tool, is released by University of Nevada.

What am I excited about? The design of future things. Where we build smart things, dumb things, pretty things, useful things, emotional things—a toaster that takes pride in the toast it makes, so it gets better and better. *Donald Norman, Neilson Norman Group Co-Founder, on innovation.*

Which ones have failed because of design flaws?

Design is never the single source of failure. The first spreadsheet, VisiCalc, was very difficult to use, but it won because it provided functionality that nothing else offered. The first really easy-to-use computer, the Xerox Star, failed because although it was very easy to use, it was expensive and it really didn't provide the functionality that was needed.

Products have multiple dimensions: fitting a real need, reliability, aesthetics, usability, price. All have to be correct for the product to succeed. No single factor dominates. Any single factor can kill the success.

What does the future of human-centric design look like?

What am I excited about? The design of future things. Where we build smart things, dumb things, pretty things, useful things, emotional things—a toaster that takes pride in the toast it makes, so it gets better and better and better.

Mark Uland

CTO and Founder, Elsinore Technologies

You are, in a manner of speaking, a bug expert. Has the Internet caused an increase or a decrease in the number of bugs shipped?

I think Internet applications have reduced the number of bugs shipped for the following reasons: First, Web interfaces are fairly simplistic. There's only so much you can do to hurt a browser. Second, the server-centric nature of a Web application helps focus the development onto a narrower target. For example, a desktop application must handle multiple OSs with multiple versions of the same technologies. A Web-centric solution can focus on just one platform with one set of technologies. Third, the ease of publication

takes out a huge portion of the testing cycle, which is the distribution and installation of the product to the internal environment. One person installs the product on the server that will be tested and makes it available to everyone.

One key thing is that a bug's impact in a Web application more quickly effects the entire organization. Desktop applications will exhibit individual manifestations of bugs, whereas Web apps will present problems for everyone simultaneously. The upside is that fixes are faster to distribution.

What's next for team collaboration? Will we ever be able to have global work-at-home teams?

All of the basic tools for successful team collaboration are out there. We can communicate extremely efficiently through email, Webex, and instant messaging—all of those things we used to need to hang around the office for. We can also share and control the intellectual property—such as the project plan, source code, and defect management—that used to require us to be bolted to our internal network. What's next is for organizations to feel comfortable using them for larger and larger projects.

More and more companies will be making use of work-at-home teams as a form of compensation, exactly the same as relaxed attire was used in the past.

Do you think services will become a big market for companies like yours that now sell software?

We believe that Web service technology such as .Net will become key to our customers. They want integrated incident management in all phases of their product and project cycle. We provide access to incident management data through the Web, desktop, office products, email clients, and developer tools. This is difficult to maintain across

- InterNIC created by NSF.
- Mosaic browser released by NCSA and Marc Andreessen.

1993

1994

- HotWired displays first banner ads (for Zima and AT&T).
- Netscape 1.0 released.

1995

- Microsoft releases Internet Explorer 1.0.
- RealAudio becomes the first audio streaming package.
- Sun launches Java.
- Apache 1.0 released.
- Netscape and Sun introduce JavaScript.

1996

- W3C releases CSS1 recommendation.
- W3C releases XML working draft.
- Macromedia releases Flash.

all of the operating systems and desktop technologies. A Web service will reduce the complexity of providing these services, while simplifying stuff too.

On the other hand, a product where you maintain and communicate intellectual property, such as defect management or source code control, is an unlikely candidate for a Web service. Customers are very unlikely to want to put their information into someone else's hands. You can see the pushback with the Microsoft-hosted Web service initiatives. The customers really like what the services would do, but they want them in their environment and are very uncomfortable with someone else holding the data. I don't think that's anti-Microsoft behavior. It's a natural desire to control one's own intellectual property.



Jesse James Garrett
Partner, Adaptive Path

Who on the Web is doing something unique that inspires you?

I'm immersed in writing my book at the moment, so what I find inspirational is great content. Joel Spolsky's essays on software development (www.joelonsoftware.com) are like nothing else out there: invariably smart, insightful, and highly readable.

What new tools will change how information architects work?

I think we've only just begun to scratch the surface of what can be done with dynamic architectures in content management systems (CMSs). The transition from working with static architectures is such a huge leap forward that we haven't yet worked out how to approach it. It's as if we've just learned how to ride a bicycle and now we're sitting at the controls of a 747. We're getting pretty good at taxiing around, but we're not off the ground yet.

How do you see site architecture evolving over the next few years?

I think we'll see the formalization of a lot of the architectural conventions that have begun to evolve over the last few years. A lot of navigational conventions have come and gone—such as 1997's "Yellow Fever" spate of left-hand nav bars, or the avalanche of Amazon-style tabs we've seen more recently—but architectural conventions seem to take more time to develop. But I would not be surprised if, for example, five years from now you could find press

releases given exactly the same architectural treatment on every Fortune 500 site.

In terms of structure, what sites do you admire?

All Music Guide offers lots of ways to navigate through massive amounts of content without overwhelming the user; now if only they would get rid of those inscrutable URLs. Salon seems to be setting the standard among daily publications for keeping archived content readily accessible. I'm fascinated by the Apple site—[the designers] defy a lot of the conventional wisdom about IA, and yet the site works just fine.

What's the most interesting debate raging in the IA field?

The big debate right now is the role of research in the work we do. One school of thought holds that IA should be a research-driven field, with architectural decisions derived directly from the results of research studies. The other viewpoint, which I favor, is that research should inform, but not dictate our work, and that our discipline should focus on creative skills development rather than research methodologies.



Evan Williams
CEO, Pyra

What technology is on the horizon that will change how you work?

I'm optimistic that better collaboration tools are around the corner. Collaboration amongst creative people is something I've been intensely interested in for years, and I never talk to anyone who isn't constantly frustrated by their information-management/project-tracking tools.

I'm impressed with Groove (though not enough to use it), and Six Degrees (www.creo.com/sixdegrees/) sounds interesting.

Blogger is Pyra's best-known project. Are blogs really revolutionary, or the same old Web in a new package?

I don't think blogs are revolutionary. I think the Web is revolutionary and blogs are simply unleashing some of its latent potential. The idea that virtually anyone can go online and say what they think with very little technical expertise or out-of-pocket expense (and, can actually garner an audience) is huge.

- ICANN selects new TLDs: .aero, .biz, .coop, .info, .museum, .name, and .pro.
- SOAP specification given Note status by W3C.
- XHTML recommendation released by W3C.
- VoiceXML specification released by VoiceXML forum.

WAP specification released by WAP forum.

1998

1999

Napster makes its debut.

2000

2002

Microsoft releases .Net.

Contributing Sources: Robert Zakon, "Hobbes' Internet Timeline;" and Jeff Swenerton.

adaptive path

Join Adaptive Path partners **Jeffrey Veen**, **Peter Merholz** and **Lane Becker** for a 3-day workshop on collaborative strategies for web site design that incorporate user goals, business requirements, and organizational awareness into your product development process.

Get Hands-on Experience

In-class activities will intimately acquaint you with user experience design methods, and all attendees receive sample deliverables and editable document templates that can be used in your work environment immediately.



USER EXPERIENCE WORKSHOPS



Chicago

29-31 May 2002

Guest: **Marc Rettig**



Atlanta

10-12 June 2002

Guest: **Bill Curtis**



Washington DC

27-27 Aug 2002

Guest: **Thom Haller**



Dallas

22-24 Sept 2002

Guest: **Paula Thornton**



San Francisco

2-4 Oct 2002

Guest: **Bob Slote**

COMING TO FIVE CITIES IN 2002

Sponsors

<http://adaptivepath.com/events>



VOICES THAT MATTER™

Internet Strategies for Technology Leaders

new.architect

Exclusive Media sponsor



500watts

This democratization-of-media concept has been promised since the beginning of the Internet, so it seems almost quaint now. But, like all new mediums, it took awhile to figure out exactly what the Web was good for and how to make it work. That's why blogs didn't take off until a few years in. Now that we've reached a critical mass, blogs—and by that, I mean the ideas, the community, and the tools and services being built around them—are leading a revolution in the way many types of news and information are published and consumed.

You've done some work adding Blogger to corporate CMS systems. How do you see blogs being used in a corporate environment?

I'm convinced that blogs will become more and more prevalent on most types of Web sites, including corporate sites. The format just makes sense: the frequency, simplicity, and brevity of blog content. Done right, a blog can turn a dead site into a living entity that helps communicate over time and build relationships with whomever the intended audience is, which is usually the (often failed) point of corporate sites.

There's also the concept of blogs on the other side of the firewall: intranet blogs, which are great tools for team collaboration, knowledge management, and internal communications.

Beau Senard

Management Consultant and Certified Project Management Professional

Do you think instant messaging (IM) or other P2P applications will affect the business landscape?

Two things resist IM adoption: first, this application dredges up the all-seeing ghost of big brother; and second, users (and IT departments) must see more upside than downside.

Another inhibitor to corporate adoption, especially for B2B applications, is the lack of a single IM standard. The problem is political though. Vendors still see propriety as a competitive advantage. Such issues are surmountable, but they require significant perceived benefits for the effort to succeed. Products such as AOL Instant Messenger; AOL ICQ; Internet Relay Chat; Yahoo Messenger; and Microsoft MSN Messenger don't communicate with one another. With few exceptions, neither does so-called corporate IM software, such as Lotus' Sametime and eRoom Technology's eRoom.

Kenneth Lacy

CIO, United Parcel Service

Is UPS keeping its eye on emerging Internet technologies?

UPS has an advanced technology group that researches, prototypes, evaluates, and tracks emerging technologies—from the moment we hear about them until they become obsolete—to determine if they could be of use. We certainly see continued innovation in the Web and Internet space. For example, XML has become a standard for us.

Recently, XML usage at UPS has been growing. Our XML steering committee has been working on various efforts with participants from many of our IT groups.

How does UPS determine that a new development should become part of UPS's business infrastructure?

Considerations include: What does it bring to the business—improved services, revenue? Is it cost effective? Does the cost justify the means? Does it comply with our IT standards? Will it integrate with our IT infrastructure?

Based on studies by the UPS Advanced Technology Group, the IT Governance committee oversees the direction of UPS technology, makes sure technology is in keeping with UPS business strategy and standards, and approves new technologies or refreshment of existing technologies.

Is there any one idea that you became aware of over the years that you would have liked to see take off, but that never really seemed to catch on?

In the late 1990s, UPS tested RFID tags on packages to facilitate our internal handling and sorting processes. Although the results of the study were promising, UPS found the costs associated with development and deployment prohibitive.



Louise Kirkbride

CEO, Broad Daylight

Why do so many companies have a hard time providing good customer service?

The fundamentals of good service haven't changed. As a customer, I still want you to answer my question quickly, easily, and accurately.

What has changed is the variety of ways customers can request that answer, and the ways companies can deliver it. Service is not just in person or on the phone any more—it's on the Web, in your email, and on your wireless device.

Companies that offer customers the option of getting support through alternate channels reap an added benefit: major cost savings. According to Forrester research, answering a question over the phone now costs companies \$12, on average, versus \$6 to answer via email. Answering the same question over the Web costs just four cents. So giving customers the option of Web self-help can provide instant gratification for customers, while reducing costs for the company.

In your opinion, which Internet technology has most impacted business?

Email has always been the killer app of the Internet. It has been adopted by every demographic segment of society with gusto. Companies benefit from it in innumerable ways, but they are also grappling with how to process the flood of email they get everyday.

And which technologies could we have done without?

The technologies that try to pretend that a Web site is actually a person are pretty annoying—for example, when you're on a site and a little bot pops up to chat with you. Customers know that there are real people on the other side of a Web site. Creating phony people to interact with does not do a service to customers.

Please answer all questions, sign and date this card. Incomplete forms cannot be processed.

1 Do you wish to receive/continue to receive *New Architect* magazine?

☐ **Yes** ☐ **No**

Signature **X**
(required)

Date
(required)

Name

Company Name

Address

City State Zip

Phone

FAX

Email

By providing us your email address, you grant *New Architect* permission to contact you via e-mail concerning your magazine subscription as well as other *New Architect* products and services.

Question 4 continued...

Business Unit/Department Management

- 50 ☐ Marketing
51 ☐ Human Resources
52 ☐ Business Development/R&D

99 ☐ Other (Please specify)

5 What is the primary business of your organization? (Select only one)

General Industry

- 01 ☐ Media (Broadcast/Publishing/
Entertainment/Advertising)
02 ☐ Manufacturer (other than computer)
03 ☐ Finance/Banking/Accounting
04 ☐ Insurance/Legal Services/Real Estate
05 ☐ Government (federal/state/local/military)
06 ☐ Health Care/Pharmaceutical
07 ☐ Wholesale/Retail Trade
08 ☐ Non-Profit/Trade Association
09 ☐ Transportation
10 ☐ Education
11 ☐ Utilities
12 ☐ Hospitality/Travel
13 ☐ Aerospace
19 ☐ Other (please specify)

Internet/Web Industry

- 20 ☐ xSP/ISP/ Service Provider
21 ☐ Web Hosting Service
22 ☐ Web site Design
23 ☐ Internet/Web Consultant
24 ☐ Internet Hardware Manufacturer
25 ☐ Internet Applications Developer/Publisher
29 ☐ Other (please specify)

High Technology Industry

- 30 ☐ Software Developer/Publisher
31 ☐ Hardware Manufacturer
32 ☐ Network Equipment Manufacturer
33 ☐ Telecommunications
34 ☐ Cable/Satellite Service Provider
35 ☐ Consultant
36 ☐ VAR
37 ☐ Systems Integrator
49 ☐ Other (please specify)

6 In terms of acquiring new tools and technologies which of the following best describes: (Select only one in each column)

- | | You | Your Organization |
|--|-----------------------------|-----------------------------|
| Eager to try new tools and technologies | 01 <input type="checkbox"/> | 20 <input type="checkbox"/> |
| Typically wait for reviews or first release revisions before trying | 02 <input type="checkbox"/> | 21 <input type="checkbox"/> |
| Wait until product or services are in widespread use | 03 <input type="checkbox"/> | 22 <input type="checkbox"/> |
| Invest when we know we are falling behind | 04 <input type="checkbox"/> | 23 <input type="checkbox"/> |
| Generally not interested and will stick with what we have because it works | 05 <input type="checkbox"/> | 24 <input type="checkbox"/> |

7 Which of the following products, services and/or applications do you specify, recommend, approve, purchase or influence the purchase for your company or your clients' company? (Select all that apply)

	Currently	Within the next 12 months
Network Infrastructure		
Switches	01 <input type="checkbox"/>	10 <input type="checkbox"/>
Routers	02 <input type="checkbox"/>	11 <input type="checkbox"/>
WAN	03 <input type="checkbox"/>	12 <input type="checkbox"/>
Transmission equipment	04 <input type="checkbox"/>	13 <input type="checkbox"/>
LANs, SANs and network management	05 <input type="checkbox"/>	14 <input type="checkbox"/>
Other	09 <input type="checkbox"/>	19 <input type="checkbox"/>
Hardware		
Servers	20 <input type="checkbox"/>	40 <input type="checkbox"/>
Workstations	21 <input type="checkbox"/>	41 <input type="checkbox"/>
Personal computers	22 <input type="checkbox"/>	42 <input type="checkbox"/>
Monitors	23 <input type="checkbox"/>	43 <input type="checkbox"/>
Mobile/wireless hardware	24 <input type="checkbox"/>	44 <input type="checkbox"/>
Digital imaging devices	25 <input type="checkbox"/>	45 <input type="checkbox"/>
Internet appliances	26 <input type="checkbox"/>	46 <input type="checkbox"/>
Storage	27 <input type="checkbox"/>	47 <input type="checkbox"/>
Printers	28 <input type="checkbox"/>	48 <input type="checkbox"/>
Other	39 <input type="checkbox"/>	59 <input type="checkbox"/>
Software		
Operating systems	60 <input type="checkbox"/>	80 <input type="checkbox"/>
Database server clients	61 <input type="checkbox"/>	81 <input type="checkbox"/>
Business applications	62 <input type="checkbox"/>	82 <input type="checkbox"/>
Content development	63 <input type="checkbox"/>	83 <input type="checkbox"/>
Security/firewalls	64 <input type="checkbox"/>	84 <input type="checkbox"/>
Development tools	65 <input type="checkbox"/>	85 <input type="checkbox"/>
Web server software	66 <input type="checkbox"/>	86 <input type="checkbox"/>
CRM/eCRM,	67 <input type="checkbox"/>	87 <input type="checkbox"/>
Database Management/		
Document Retrieval	68 <input type="checkbox"/>	88 <input type="checkbox"/>
Statistical site analysis	69 <input type="checkbox"/>	89 <input type="checkbox"/>
Site management	70 <input type="checkbox"/>	90 <input type="checkbox"/>
Other	79 <input type="checkbox"/>	99 <input type="checkbox"/>
Electronic Commerce Tools		
Hardware	100 <input type="checkbox"/>	110 <input type="checkbox"/>
Software	101 <input type="checkbox"/>	111 <input type="checkbox"/>
Network infrastructure	102 <input type="checkbox"/>	112 <input type="checkbox"/>
Other	109 <input type="checkbox"/>	119 <input type="checkbox"/>
Internet Access/Web Hosting		
Internet access	120 <input type="checkbox"/>	130 <input type="checkbox"/>
ISDN services	121 <input type="checkbox"/>	131 <input type="checkbox"/>
VPN	122 <input type="checkbox"/>	132 <input type="checkbox"/>
Dedicated leased lines	123 <input type="checkbox"/>	133 <input type="checkbox"/>
Fractional T-1/T-1 Services,		
T-3 Services	124 <input type="checkbox"/>	134 <input type="checkbox"/>
Web services	125 <input type="checkbox"/>	135 <input type="checkbox"/>
Other	129 <input type="checkbox"/>	139 <input type="checkbox"/>
Application Development Tools		
Languages	140 <input type="checkbox"/>	160 <input type="checkbox"/>
IDEs	141 <input type="checkbox"/>	161 <input type="checkbox"/>
Encryption packages	142 <input type="checkbox"/>	162 <input type="checkbox"/>
HTML authoring	143 <input type="checkbox"/>	163 <input type="checkbox"/>
Compilers	144 <input type="checkbox"/>	164 <input type="checkbox"/>
Platforms	145 <input type="checkbox"/>	165 <input type="checkbox"/>
Mobile authoring	146 <input type="checkbox"/>	166 <input type="checkbox"/>
Multimedia tools	147 <input type="checkbox"/>	167 <input type="checkbox"/>
Other	159 <input type="checkbox"/>	179 <input type="checkbox"/>
Outsourcing/ Services		
Consulting	180 <input type="checkbox"/>	190 <input type="checkbox"/>
ASPs/WASPs	181 <input type="checkbox"/>	191 <input type="checkbox"/>
Web design	182 <input type="checkbox"/>	192 <input type="checkbox"/>
Web site development	183 <input type="checkbox"/>	193 <input type="checkbox"/>
Systems integration	184 <input type="checkbox"/>	194 <input type="checkbox"/>
Application hosting	185 <input type="checkbox"/>	195 <input type="checkbox"/>
Customer relationship	186 <input type="checkbox"/>	196 <input type="checkbox"/>
Other	189 <input type="checkbox"/>	199 <input type="checkbox"/>

For faster service fax both sides of this form to

847-647-5972

or visit us at
www.newarchitectmag.com/subscriber/

Continued on back. →

8 How much will your organization (or your consulting clients' combined) spend in the following categories for Internet-related products, applications, or services in the next 12 months?
(Select one amount from each category)

	Hardware	Software	Access/ Hosting Servers	Outsourcing/ Consultants
\$100,000,000 or more	01 <input type="checkbox"/>	20 <input type="checkbox"/>	30 <input type="checkbox"/>	40 <input type="checkbox"/>
\$25,000,000 - \$99,000,000	02 <input type="checkbox"/>	21 <input type="checkbox"/>	31 <input type="checkbox"/>	41 <input type="checkbox"/>
\$10,000,000 - \$24,999,999	03 <input type="checkbox"/>	22 <input type="checkbox"/>	32 <input type="checkbox"/>	42 <input type="checkbox"/>
\$5,000,000 - \$9,999,999	04 <input type="checkbox"/>	23 <input type="checkbox"/>	33 <input type="checkbox"/>	43 <input type="checkbox"/>
\$1,000,000 - \$4,999,999	05 <input type="checkbox"/>	24 <input type="checkbox"/>	34 <input type="checkbox"/>	44 <input type="checkbox"/>
\$500,000 - \$999,999	06 <input type="checkbox"/>	25 <input type="checkbox"/>	35 <input type="checkbox"/>	45 <input type="checkbox"/>
\$100,000 - \$499,999	07 <input type="checkbox"/>	26 <input type="checkbox"/>	36 <input type="checkbox"/>	46 <input type="checkbox"/>
Less than \$100,000	08 <input type="checkbox"/>	27 <input type="checkbox"/>	37 <input type="checkbox"/>	47 <input type="checkbox"/>

9 Approximately how many people are employed in your entire organization (or your consulting clients' combined)?

- 01 ☐ More than 20,000 Employees
02 ☐ 10,000 - 19,999 Employees
03 ☐ 5,000 - 9,999 Employees
04 ☐ 1,000 - 4,999 Employees
05 ☐ 500 - 999 Employees
06 ☐ 100 - 499 Employees
07 ☐ Less than 100 Employees

10 What is the highest level for which you determine the need, evaluate, recommend, specify, buy or approve hardware, software or access services for your organization's (or your consulting clients') Internet initiatives?
(Select only one)

- 01 ☐ Several Different Companies
02 ☐ Entire Company
03 ☐ Company-wide
04 ☐ Several Divisions or Sites
05 ☐ Several Departments at your location
06 ☐ Department/Workgroup
07 ☐ Myself

11 Please check the boxes below to let us know the types of information that may be of interest to you and that we may send you via e-mail:
(Select all that apply)

- 01 ☐ Product and service offerings from other CMP Media LLC properties.
02 ☐ Product and service offering from other relevant businesses and organizations.

Publisher reserves the right to determine qualification for free subscriptions.


▼ PLEASE FOLD ALONG THIS LINE ▼



SKOKIE IL 60076-9483

PO BOX 1246

new architect
Internet Strategies for Technology Leaders



POSTAGE WILL BE PAID BY ADDRESSEE

BUSINESS REPLY MAIL
FIRST-CLASS MAIL PERMIT NO 1099 SKOKIE IL

NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES



PLEASE SEAL WITH TAPE ALONG THIS EDGE ONLY, THIS PANEL OUT. THANK YOU.

It's as if we've just learned how to ride a bicycle and now we're sitting at the controls of a 747. We're getting pretty good at taxiing around, but we're not off the ground yet.

Jesse James Garrett, Adaptive Path Partner, on information architecture.

CRM has been criticized as a technology that is more trouble and cost than it's worth. How do you respond to the critics?

Well, a few years ago I read a story that claimed that you really could not quantify any increase in productivity due to desktop automation. If so, I think it was a failure in the analysis, not the technology. I believe the same is true in this case. [CRM issues] are really hard problems to solve, and each company handles its customer relationships differently.

The companies most likely to succeed with CRM technology are those that focus on solving a specific business problem, and recognize that human behavior does not change overnight.

What does the future of customer service look like?

Customers will find more and more ways to contact companies, and they will do so 24 hours a day, 7 days a week. Questions about Christmas gifts will have to be answered Christmas day. Customers will want to be able to conduct business whenever they want, wherever they want.

Kevin Lynch

Chief Software Architect, Macromedia

What technology is on the horizon that will change how you work?

Rich clients. There's a fundamental shift coming in how applications are built and delivered over the Internet that will work across personal computers and devices, function both online and offline, and seamlessly blend content, application UI, and communications. Rich clients combined with the new 2.5G wireless networks that are becoming available soon will, I believe, mark a significant shift in enabling mobile wireless applications to finally become truly useful and liberating.

Which companies in our industry have done the best job of keeping pace with market demands?

Essentially, the ones that have remained in business! The most successful companies have made their products extensible to enable an ecosystem of developers to innovate and build on them. This also results in those products keeping pace with market demand and technological shifts through quick development by the community itself. This is definitely something that we did early on with our products, and clearly what Microsoft, IBM, Sun, BEA, and other systems companies have pursued.

How have notions of usability changed how your developers build?

We have always included usability testing in our process of develop-

ing applications. We strive to make them both very approachable to new users while also being extremely productive for experienced users, which is a challenging balance. One of the main reasons the Web itself was able to grow so quickly was the simplicity of clicking on links to browse documents.

As HTML has grown, and sophistication of Web page design has developed, it's requiring more skill and experience to maintain the simplicity of the original Web. We have been focusing on how to help designers and developers using our software produce more effective, usable results. For example, Dreamweaver now supports checking a page through UsableNet, which gathers the common knowledge on usable HTML design and can automatically flag common issues of usability and accessibility. As people start to use the Web for applications, we're seeing that HTML's page-based metaphor isn't quite as strong a way to deliver the most effective application user interfaces. For some applications that require a more sophisticated user interface, we are seeing developers begin to take advantage of rich clients inside the Web browser. We believe there is a direct relationship between the usability of applications on the Web and the revenue and cost structure of the business. The most usable applications can increase revenues and decrease costs, and the sooner businesses realize this, the more we will see dramatically improved usability across the Web.

Can Macromedia Flash and SVG coexist?

Yes, definitely. They coexist today. Flash and SVG are actually not really that equivalent. While Flash displays vector graphics, it is much more than that today with its support for MP3 audio, Sorenson video, and rich application user interfaces connected to application servers. Flash uses a compressed binary format (SWF) to deliver its content, where SVG uses a larger text format to describe its content via XML. Also, the Flash Player is very small (less than 500K), while the Adobe SVG plugin is anywhere from 2.2MB to 4.6MB in size, depending on platform.

What's the most interesting debate raging in the design and development community right now?

Usability—balancing emotion with intellect in designing effective content. This is best exemplified by the extremes of Jakob Nielsen, the experimentation by Joshua Davis, and the blending of these in Jeffrey Veen's *The Art and Science of Web Design*.

What question would you love to answer that no journalist ever asks?

What's the next big Internet application after the Web? ■

XML-based description formats like RDF and DAML+OIL are slowly fulfilling our hopes for a more meaningful, helpful Internet.

the of the languages semantic web

by Uche Ogbuji

To create the Web as we know it, Tim Berners-Lee put aside much of the existing research on hypertext technologies and built a simple system that was easy to understand, use, and maintain. This simplification became an important factor in the Web's rapid growth. Despite this success, the realities of information management are illuminating some problems of simplification. While the Web continues to be useful for retrieving information from individuals or organizations of close collaborators, it is much harder to use if you want to gain a broad understanding of a particular subject.

For example, while we can visit the Burton snowboards Web site to find out what products the company offers, read about its corporate policies and philosophies, and even browse a selection of links about snowboarding in general, it's much more difficult to find a wider perspective on snowboarding as an industry and interest.

It's even harder to bind together the many Web sites that discuss snowboarding.

This is where the Semantic Web comes in. The Semantic Web is a vision of a next-generation network that lets content publishers provide notations designed to express a crude "meaning" of the page, instead of merely dumping arbitrary text onto a page. Autonomous agent software can then use this information to organize and filter data to meet the user's needs.

There has been much effort to refactor the Web more along these lines since the success of the current Web. Proponents of this goal often refer to it as the Intelligent Web. For those who focus on the problem of how to express the context—or, the semantics—of content in distributed systems like the Web, this goal is called the Semantic Web.

Even though this next-generation Web has yet to become a reality, much of the current work on the Semantic Web centers on a variety of

technologies that are already in widespread, practical use. In particular, the Resource Description Framework (RDF)—which lets content creators express structured metadata statements describing URIs.

limits of today's web

With the current state of the Web, there are only two real methods of gaining broader information about documents. The first is to use a directory or portal site, and thus rely on human editors to scour the Web and appropriately categorize pages and their associated links.

Such portals are the heroes of today's Web. After all, the most effective information management tool on Earth is still the human librarian, and probably will be for years to come. The problem is that directories take tremendous effort to maintain. Finding new links, updating old ones, and maintaining the database technology add to a portal's administrative burden and operating costs.

Search engines are the alternative. Good search engines pay special attention to metadata in the pages that they spider and add to their index databases. In the simplest case, this metadata might take the form of content in `<meta>` tags. More advanced search engines, like Google, rely on more subtle information. For instance, Google's widely touted algorithm evaluates not only the occurrence of keywords on a page, but also the number of outside links to the page itself, as a measure of its importance or popularity.

Search engines take less human effort on the content management end, but they require a frightfully large resource investment. It's also very difficult to produce valuable indices efficiently. It's no secret that some of the most advanced search engines are so primitive that queries often turn up an unmanageable number of poorly differentiated hits. A user who tries to finely craft his or her search to zero in on a point risks filtering out potentially relevant search results.

The Web needs to support something in between portals and search engines. Of course, until there's a server as sophisticated as HAL 9000 (but, hopefully, not as neurotic), we probably won't be able to completely replace the human portal editor with a computer program. But if we could provide standardized means for Web publishers to catalog and classify their own content, then we could develop more effective agents that work on this substrate of better-organized information.

The result of having better standard metadata would be a Web where users and agents could directly tap the latent information in linked and related pages. This would help free us from having to scour for information site by site, and from relying on portals and search engines. It wouldn't be hard to outfit each user with personal portal generators and search agents tailored to their particular interests, needs, and constraints. These agents might even be configured to learn and respond to personal details with the help of artificial intelligence techniques.

the semantic web's challenges

It's fine to talk about enabling each Web publisher to properly place content in context, but there are several problems to overcome before any such initiative will gain critical mass:

- **Complexity.** Any technology that the average Web developer can't grasp in a day and apply proficiently in a week is doomed. In addition, a successful technology will have to be integrated into current Web development and maintenance tools. Semantics are quite arcane, and it won't be easy for semantic technologies to meet this criteria.

- **Abuse.** Practices like meta-tag spamming, and even trademark hijacking, show that any system that lets people set their own context is subject to abuse. Knowing the value of the Burton snowboards brand, another unscrupulous manufacturer might want to tell an agent that it is the Burton company in hopes of directing some undeserved attention to its site. Semantic Web technologies will need a mostly automated system for establishing trust in the assertions that Web publishers make. This concept is often referred to as the *Web of trust*.

- **Proprietary Technology.** Because of the diversity in developers and development tools, Semantic Web technology will have to be politically and technically open for implementation and use. If it requires royalty payments to any party, open source advocates and competing Web technology vendors will boycott it. If it requires a specific plug-in or module, most developers and users won't even bother installing it.

Semantic Web proponents are looking to XML and RDF to meet these challenges. XML would let a publisher use markup that differentiates a catalog entry of a snowboard product from an independent review of the same item. However, this method relies on custom tags, and agents need a way to grasp the "meaning" in such tags—a facility called *semantic transparency*. Web metadata is the key to providing it. Because of its importance, the W3C developed RDF as a standard for Web metadata.

inside rdf

RDF is indeed quite simple at its core, though it can get hairy in short order. It is a model of statements made about resources. A resource is anything with an associated URI. In practice, it's most often a document on the Web, but it can be anything to which people have agreed to assign a URI. In this way, one could even use RDF to make statements about abstractions like peace, or even imaginary entities like Gandalf the Wizard. RDF's statements are hardly as complex as those we use in natural language. They have a uniform structure of three parts: predicate, subject, and object. For example: The author [predicate] of *The Lord of the Rings* [subject] is J.R.R. Tolkien [object].

This simplicity and uniformity make RDF's statements generic. They can be used to encode the above natural-language statement, as well as, say, an object-oriented model. For example, if you had written a class called `Person`, and that class was instantiated as an object called `jrrt`, your statements would be: `Person's type is Class`; and `jrrt's type is Person`. The connections between various subjects and objects can be much more complex than this, of course, especially if you think about inheritance and properties and other attributes that classes take on. As you might guess, this approach can be tedious if done by hand.

RDF lets you express such statements in a formal way that software agents can read and act on. It lets us express a collection of statements as a graph, as a series of (subject, predicate, object) triples, or even in XML form. The first form is the most convenient for communication between people, the second for efficient processing, and the third for flexible communication with agent software.

If a portal were to create a directory of snowboarding sites, it could use such an RDF/XML document to help RDF-enabled agents and tools better understand the information that the sites offer. Example 1 is loosely based on the format used by the Open Directory Project (www.dmoz.org), a community effort to build a universal Web site directory.

EXAMPLE 1: AN RDF DIRECTORY

```
<?xml version="1.0"?>
<rdf:RDF xmlns:rdf="http://www.w3.org/1999/02/
  22-rdf-syntax-ns#"
  xmlns:s="http://rdfinference.org/eg/
    snowboard/metadata/"
  xmlns:dc="http://purl.org/dc/
    elements/1.1/">
  <rdf:Description rdf:about=
    "http://rdfinference.org/eg/snowboard/
      metadata/topics">
    <dc:title>Snowboarding</dc:title>
    <s:SubTopic rdf:resource=
      "http://rdfinference.org/eg/snowboard/
        metadata/topics/references"/>
  </rdf:Description>
  <rdf:Description rdf:about=
    "http://rdfinference.org/eg/snowboard/
      metadata/topics/references">
    <dc:title>Snowboarding
      references</dc:title>
    <s:Link>
      <rdf:Description
        rdf:about="http://www.geocities.com/
          iliktoast/">
        <dc:title>The Snowboard
          Encyclopedia</dc:title>
        <dc:description>Provides info on
          gear, history, etc...
        </dc:description>
      </rdf:Description>
    </s:Link>
  </rdf:Description>
```

A group of related sites can be marked up in an RDF document in a way that is clear to autonomous agents.

The first document element, `rdf:RDF`, tells an RDF parser that the child elements can be interpreted as RDF constructs. Note the namespace declarations. The first one defines the core namespace for RDF constructs. The second is a special namespace that is controlled by our fictitious snowboarding Webmaster community.

The first child is an `rdf:Description` element, which tells the RDF parser that we have a resource to describe. The `rdf:about` attribute notes the URI of the resource being described. In this case, it's a local resource to the community site that refers to the overall topic of snowboarding (according to some agreement the community will have made).

The `dc:title` child is known as a property element, and gives the predicate of a statement to be made. The object of the statement is given by the text content. In this case, we are making a statement similar to this: Snowboarding is the title of the topic identified at `http://rdfinference.org/eg/snowboard/metadata/topics/references`.

Next comes another property element, `s:SubTopic`, with another resource provided as an object rather than a text string. This one relates the topic to its sub-topics, establishing the directory hierarchy. The next description is for this sub-topic. It also has a title and a relationship to a link, using the `s:Link` that is specified as another resource—which is a child of the property element—and thus, becomes the object of the link statement. In this case, the link is to the site `www.geocities.com/iliktoast/`, which we further describe with a title and description as well, also defined according to DCMI.

Of course, it would be best for the community if each Web page could maintain its own metadata. The RDF specification provides a convention for people to place RDF within HTML pages. Example 2 illustrates how the maintainers of various snowboarding sites might use RDF to do this for their own pages.

The empty `rdf:about=""` attribute is a special URI convention that refers to the current document. Other than that, the code is similar to that in the RDF directory. Note that this data can be maintained in tandem with regular HTML `<meta>` tags to support existing search engines and RDF agents. One hopes that vendors of popular Web authoring tools will soon produce products that automatically represent metadata in both RDF and `<meta>` formats.

the scheme of things

RDF itself is a handy way to describe resources. Widespread use of such a facility could alleviate many of the current problems with the Web. But RDF by itself only gets us part way toward realizing the Semantic Web, in which agents can infer relationships and act on them.

Classification is extremely important on the Semantic Web. Each community of related interests defines categories for the matters that it discusses. For instance, the snowboarding community defines items such as snowboards, parks, tricks, and manufacturers. The definition of a manufacturer in snowboarding terms is related to the definition of a manufacturer in the general business sense. The snowboarding community can enshrine these definitions by creating a schema for its RDF models. The W3C is standardizing a simple mechanism for RDF schemas (W3C RDFS). If you're interested in seeing what the actual code looks like, I've provided Example 3 online at `www.newarchitectmag.com`.

W3C RDFS is itself expressed in RDF format. The first stanza of an RDFS document, like that in Example 3, describes a class of item that we identify with the given URI and label, for instance, Snowboard. The comment is just a useful documentation item. The second stanza might describe a class we label Snowboard Manufacturer. We subclass this label from the more general concept of *manufacturer* that's defined in the RosettaNet general business dictionary. RosettaNet is an organization for standardizing business-to-business communication using computers. Perhaps from this connection, you can understand how an agent designed for processing information on general business matters could gain at least a foothold of understanding if it were to come across a snowboarding site. This is an essential trick of the Semantic Web.

We use the same trick to define a rider as a specialization of *person* defined in FOAF, a well-known schema for personal and organizational contact information. We also define a couple of properties in the schema—resources that can be used as the predicates of RDF statements. W3C RDFS lets us apply some simple constraints on properties. For instance, `rdf:domain` only lets a property be used on a certain class of resources, `rdf:range` declares that the value of the prop-

erty must be of a certain class of resources. So, for example, we say that only a Rider can have an Endorsement property, and that the value of all such properties must be Snowboards.

With this schema in place, the snowboarding community would have a formal basis for saying things like "Chris Englesmann endorses the Fatbob snowboard, manufactured by K2." Of course, this isn't to say that all of the content on an RDF-enabled site would be translated to abstract graphs, or long-winded XML representations thereof. RDF would rather be used in content headers—in the HTML <head> section, for example—to make formal statements about the content that guides an agent in placing it in context.

nebulous knowledge

Schemas take us a step toward the Semantic Web, but not all the way. In the heyday of artificial intelligence, scientists were puzzled by a crucial point. Although computers were beginning to overtake the human brain in terms of sheer processing speed and storage capacity, they still didn't compare to human intelligence. At least one reason for this is that the brain doesn't stubbornly store and categorize every scrap of every detail that we use as the basis of thought. The brain is a miracle because it can make connections between partially-stored information, and assemble this into intelligence when necessary. To achieve this level of understanding with RDF and RDFS, countless resources on servers all over the world would have to be methodically classified and described. This, of course, is completely unrealistic.

The Semantic Web won't be possible until agents have the means to figure out some things by themselves, given the data they have to work with. Fortunately, artificial intelligence gives us two tools to help make this possible. First, knowledge representation is a field that defines how we might represent, in computers, some of what is stored between our ears. This would give computers a fighting chance at synthesizing unclassified data at a useful speed. Second, inference is a way of using formal logic to approximate further knowledge from that which is already known. All of this forms a system of representing and synthesizing knowledge that is often referred to as an ontology.

The leading ontology system for RDF is the DARPA Agent Markup Language (DAML). DARPA, for those who may have forgotten, is the group that brought us the Internet itself. DAML incorporated useful concepts from the Ontology Inference Layer (OIL), a European project to provide some AI primitives in RDF form. The resulting language is DAML+OIL. (Visit www.daml.org for more information.)

DAML+OIL lets us formally express ontologies. W3C RDFS provides primitive classification and simple rules for this, but DAML+OIL goes much further. For instance, DAML+OIL can express that "any snowboard with plate bindings is a race board," which makes it unnecessary to then explicitly flag every race board. You might see in this some of the flavor of business rules, which are known in software development circles as the programmatic expression of mandates for the way data must be processed. In fact, one way to look at DAML+OIL is as the business rules for the Semantic Web, yet it's much more flexible than most business-rules-languages in common use.

Most of DAML+OIL's power comes from primitives for expressing classifications, as the race boards example illustrates. DAML+OIL provides a toolbox of class expressions, which bring the power of

EXAMPLE 2: EMBEDDING RDF IN HTML

```
<rdf:RDF xmlns:rdf="http://www.w3.org/1999/02/
22-rdf-syntax-ns#"
xmlns:dc="http://purl.org/dc/elements/1.1/">
  <rdf:Description rdf:about="">
    <dc:title>The Snowboard
    Encyclopedia</dc:title>
    <dc:description>Provides information on
    gear, history, etc...
    </dc:description>
  </rdf:Description>
</rdf:RDF>
```

Web publishers can embed advanced metadata in their pages by placing RDF data like this in the <head> of their HTML documents.

mathematical logic and set theory to the tricky and important task of mapping ontologies through classifications.

miles to go before we sleep

The Semantic Web is still a way off, if it's attainable at all. To date, RDF and DAML+OIL are our best efforts at reaching it. They address a good number of the problems with the present state of the Web, and further enhancements are on the way. For example, a system of statements that's managed at a certification authority could help establish the validity of RDF statements to minimize metadata spam and other security problems.

There are already automated tools to help generate RDF for existing Web pages, which should aid migration. In fact, RDF's great strength is that even without the Semantic Web in place, it has proven to be a practical and usable technology in assorted areas of computing. It can be used for data descriptions in highly generic and extensible databases, or for sophisticated modeling in application development. As a result, an impressive and growing selection of tools and vocabularies are defined in RDF for ready use. The integration with Web publishing tools is improving with developments like the PRISM standard for content syndication metadata. This standard is RDF-based and endorsed by an impressive cross-section of the electronic publishing and publishing tools industries.

Many of the practicalities of the Semantic Web depend on how we define victory. If it's enough to make the Web a better-managed resource for well-defined and somewhat more organized communities, then RDF and the higher-level technologies I've discussed already provide the means to do so now; it's just a matter of obtaining a critical mass of users. RDF success stories like the RDF Site Summary (RSS) for content syndication and the Musicbrainz system (www.musicbrainz.org) for digital music metadata could be the catalysts of significant progress. And soon the Web might have an alternative between the poor scalability of traditional librarians and the ineffective results of today's search engines. ■


Uche is a consultant and co-founder of Fourthought, a consulting firm that specializes in XML solutions for enterprise knowledge management applications. Contact him at uche.ogbuji@fourthought.com.

As server infrastructures become more complex, they increasingly require human intervention to keep running. IBM and other vendors believe they have the answer: let the server manage itself.

by Jay Lyman

server, heal thymself

THE QUEST FOR
AUTONOMIC
COMPUTING



Modern Web infrastructures are an information-management dream come true. Computers can manage and automate tasks that would have taken countless man-hours in the past. Ironically, however, as organizations hand increasingly greater responsibilities to their servers, the level of skill required to maintain those systems has sharply increased.

"Infrastructures are becoming more complex," says Frank Kettenstock, vice president of product marketing for Marlboro, MA-based Concord Communications, "but the average skill set of the people managing them is getting lower and lower. The people just aren't there." Kettenstock explains that the slumping economy has only made matters worse. "The need hasn't gone away, the dollars to pay for it have," he says.

The answer, according to Kettenstock, is automation and autonomy of hardware and software. To that end, Concord Communications—like IBM and several other infrastructure and software vendors—creates and markets software to automate technology management. The ultimate goal is to create what some have begun calling "self-managing, self-healing" servers.

autonomic advances

The aim of management automation software is to maximize physical resources and decrease the amount of human interaction necessary to keep IT running smoothly. Backers of self-managing, self-healing, self-configuring, self-optimizing machines—what IBM calls *autonomic* servers and software—say the increased complexity of applications and computer networks requires smarter technology for administrators.

Already, servers and software are capable of doing much more on their own than they were ten years ago. Autonomic computing ideas have manifested themselves in servers that find their own computing resources on the fly, remain up and running through glitches, and safeguard against attacks and crash. These ideas are also evident in software capable of automating database tasks, memory allocation, and so on.

The infrastructure players who are selling self-management features in their servers and software insist that these ideas are creeping into all levels of computing, from the high-end down to low-end and midrange servers. In fact, the fight for improved autonomy may eventually replace the competition for increased processing power, memory, and speed as the primary selling point for new systems. While some hurdles remain, the race to automate is on.

the eliza vision

Perhaps the loudest proponent of self-managing features is IBM, which in early 2001 announced its Project eLiza, an ongoing effort to let technology, not people, manage computing infrastructures. Already, eLiza has gained high-profile customers including Merrill Lynch and Terra Lycos, and partnerships with technology vendors like BMC Software, Candle Corporation, and Nortel Networks.

Big Blue likens the idea of autonomic servers to the human nervous system, which is capable of monitoring and adjusting heart and breathing rates without conscious intervention. The autonomic system will learn how to tweak and turn proverbial knobs, just as humans do. The goal is that eventually, servers will be able to do this without human interaction of any kind.

Some areas where IBM has been concentrating its research include: *configuration*, where servers, software, and features can add themselves to an enterprise infrastructure; *healing*, fault recovery and availability in the face of component failures; *protection* and *security*, which include access management, intrusion detection, secured backup, and real-time reporting; and *self-optimization* to maximize resources without IT overtime.

In its "autonomic computing manifesto" (www.research.ibm.com/autonomic/manifesto/), IBM stakes the future of computing on autonomy, urging the industry to take on what it calls "the Grand Challenge" of self-managing servers. "Too few of us are focused on solving...a problem springing from the very core of our success," the manifesto reads.

IBM eLiza Technical Strategist Tom Monza stresses Big Blue's extension of autonomy beyond the single-system environment. The company's autonomic computing initiative focuses on an always-on flexible system that can sift data using a platform or device-agnostic approach without the need for human intervention.

"It is a grand challenge," Monza admits. "It's not an easy problem, and we acknowledge it's something we can't do on our own."

self-management everywhere

Other hardware and software players say they're doing their part. Indeed, a survey of the computing industry reveals that self-managing technologies are beginning to appear in all strata of the marketplace. Hewlett-Packard, for example, announced its rp7410 server earlier this year, which includes a "self-healing cache," a technology that's designed to overcome CPU failure by maintaining a reserved "clean cache" space within the chip itself.

Microsoft touts its Windows 2000 and XP family of workstation/server products and SQL servers as examples of its "pioneering" role in creating easier-to-use servers, and cites its self-tuning database system project, known as AutoAdmin, as its latest thrust toward automation.

"This project has led the industry by innovating techniques to automatically pick indices and materialized views adaptively based on workload analysis," says Microsoft researcher Surajit Chaudhuri. "It was flattering to see IBM Almaden and other research organizations follow our lead," he adds.

According to Robert Shimp, Oracle vice president of database product marketing, what IBM calls its cutting edge self-management technologies have been part of Oracle's focus for some time. "All of their features are things that have been in Oracle for many years," Shimp says. "It's stuff we have today or are releasing."

Oracle's "unbreakable" marketing campaign for its 9i database translates to availability regardless of server, storage, site, or human error, according to the software maker. Down the road, Oracle will fully automate the process of updating patches and continue to make storage and space management automatic, according to Shimp. "That's a huge part of the [database administrator's] job," he says.

reducing people power

The motivation for self-managing technology, whether it's built into servers, databases, or other software, remains the same. Microsoft's Chaudhuri explains that the self-managing features of Microsoft SQL Server—including advanced memory allocation and the AutoAdmin self-tuning database system—are examples of the software giant's efforts to

first soothe database administrators (DBAs), and then supplant them. "Our goal is to make DBAs sleep easy for now and, over time, make them irrelevant and fewer in number," Chaudhuri says.

Yet many companies insist the idea is to optimize humans, not eliminate them. Oracle's Shimp points out that the administrators will still be needed. "All of these things reduce the individual tasks of what DBAs need to do," he says. "You don't want DBAs spending time on lower level services. You want them focused on the larger infrastructure."

Citing labor cost as the driving force behind automation, Shimp says database tasks such as sizing, caching, and CPU management have already been included in Oracle products. Further automation is coming in the areas of simpler configuration and updating, patch management, storage and space management, and backup and recovery.

Referring to HP's instant capacity on demand (iCOD) and workload management (WLM) products as "goal-based software," HP Product Manager Sanjiv Patel says administrators will definitely benefit from automation technology like the self-healing cache used in HP's servers. "This is something that's going to make the administrator's job easier," he says. "It's definitely a feature they'll appreciate."

present-day solutions

Whether or not autonomic controls are appreciated today, Big Blue predicts that in the future, they will be essential. IBM's Monza says eLiza and the self-management features of the project will be increasingly necessary as other computing resources improve. "As the technology becomes more affordable, the applications are becoming increasingly difficult to manage," Monza says.

IBM will soon roll out multi-system management for heterogeneous networks. The management will be aimed at IT outfits with twenty or more servers and multiple operating systems. Monza says that smaller customers can find automation at the single server level, but he admits that "the cross-system stuff is where it starts to get interesting."

Hewlett-Packard's Patel says self-management in HP's servers is most robust in higher-end products, but is showing up increasingly in mid-range servers like the rp7410.

Automation is also being featured more often in databases and application software. Microsoft's goal of replacing DBAs with self-management is one example. Meanwhile, rival Oracle claims it has been progressing since it released version 6 of its database some 12 years ago. "We've automated quite a number of database tasks, making them self-managing," says Robert Shimp.

"Essentially, there are basic policy rules that you want to be able to program into a database," Shimp says. "It's collaborative management. You have detailed information while automating low-level decision making, so DBAs aren't spending time with low-level services. You want them to be focused on the larger infrastructure and eliminate a lot of that lower-level drudge work."

dubious efforts

Skeptics say the new self-management features touted by hardware and software makers are just new names for the same old features of fault tolerance and recovery that have been incorporated into computing for decades.

Whether or not autonomic controls are appreciated today, Big Blue predicts that in the future, they will be essential.

Forrester Principal Analyst David Truog says the real issue surrounding self-management, and particularly self-healing, is whatever makes the server or software ill in the first place. "I think fundamentally what's behind these efforts is the issue of software quality and infrastructure quality—bad quality," he says. "I like the idea of self-healing software, but I think that writing better software in the first place is more effective."

Truog scoffs at the idea that much of the self-managing and healing technologies out there are anything new. "Fundamentally, this is an idea with a long history," he says. "As we become more networked, we depend more and more on technology and software." He says the idea of self-healing technology has already been applied today in many areas, including e-commerce shopping carts that are mirrored on server sites to prevent loss of the cart.

Truog does not doubt the need for self-healing servers and software in the face of increased complexity, but still feels that the large infrastructure players are doing more spinning than automating. "They've all been doing it for years," he says. "It's just that now they've put a marketing label on it."

Hewlett-Packard's Patel insists that autonomic features will continue to increase, regardless of their novelty. "We are definitely talking about manageability and availability technologies getting stronger," Patel says. "It's nothing out of left field...The technologies are making things easier for systems administrators."

Concord's Kettenstock, however, concedes that there are limitations to self-management. Concord's automation software, which uses scripted fixes to common problems, works with well-defined problems. "You tell us what to look for and we'll either tell you when it happens or go with an automated fix," Kettenstock explains. "There are limited ways in which it can be used and it's primarily for well-known problems with well-known solutions."

added complexity?

Forrester's Truog raises another point. If self-management is adding another network layer to manage, does that truly simplify a network infrastructure, or does it make it more complex? Truog feels that attention to other areas, such as software quality and layering, might be more effective at simplifying and increasing productivity.

"When you are adding additional complexity, you're adding additional points of failure," Truog says. "There's a little bit of a risk in some of these technologies and ideas." Truog agrees that as computer storage, processing, and bandwidth all increase, complexity and the need for self-management increases. But he doesn't believe autonomic technologies are necessarily going to do the job.

"All these issues are about spending money on machines to do more, so humans can be more productive," Truog says. "I'm frankly a little skeptical because of the fundamental complexity issue."

Kettenstock says another big issue with automation is gaining the trust of the people at the controls. "The biggest hurdle is the belief that [the technology] works," he says. "It looks great and it sounds like a good approach, but for the people whose necks are on the line, they get very nervous from black boxes. There's a trust factor here. You don't get productivity gains unless these people trust the automation. Gaining trust takes a long time; it takes one mistake to lose it."

Sanjiv Patel points to Hewlett-Packard's reputation as reassurance enough for most of its customers. He doubts administrators will cringe when they're cut off from busy work, and notes that IT pros will probably like the continued push toward autonomy in technology. "In the case of this type of technology, the idea is not necessarily to involve the administrator," Patel says. "It's to make it work whether a human being is there or not."

toward the future

For Concord, the next step after automation is *embedded intelligence*, or automation that closes the complexity/skill set gap. Oracle's strategy is similar. "You need to be able to put that kind of intelligence into database systems," Shrimp says, referring to such automation as regulation of system resources and access priority.

Microsoft is strongly focused on further developing the self-management framework of SQL Server and AutoAdmin. "We are building a server-based technique to automatically detect anomalous behavior in SQL Server and take corrective actions," explains

Chaudhuri. "This effort exploits our skills in real-time data analysis and mining."

HP, on the other hand, is continually working in different areas, such as operating system availability and manageability. But the goal is the same—to make technology management as easy and seamless as possible.

To date, IBM remains the most committed to incorporating automation and self-management into its server product line. The company says the benefits will include scalable storage and processing power for multiple autonomic systems, transparent routing and formatting for variable devices, better chip memory, improved network-monitoring, and smarter microprocessors that can detect errors and anticipate failures.

While other companies offer diverse concepts of what self-management truly means, IBM believes that only an industry-wide effort can truly succeed in raising computing to the next level. To that end, IBM has called for the development of the "necessary standards and open interfaces" to make automation a reality. "It's a journey," says Monza. "And who knows when we'll reach the utopian world of totally self-managing systems. This is a big effort, not just an IBM effort. It's a multi-year project." ■

Jay is a reporter for a daily tech news service. You can contact him at jaylyman@home.com.

deus ex machina

How will machines replace administrators? With artificial intelligence (AI), we're told. However, AI is a much broader concept than you might expect. Imagine how you'd feel if you asked someone how they implemented an e-commerce site, and they replied, "Networking, of course!" The reality is that AI is a general term for many specific techniques. In the quest for homeostatic servers, two AI technologies seem to be the most promising: expert systems and neural networks.

In a traditional expert system, the program processes rules while trying to decrease the number of rules examined. For example, suppose one rule is: "If pinging the server on route A has no effect, try route B." If this rule matches—that is, if pinging the server on route A has no effect—then there is no reason to try other rules that require connectivity via route A. This type of problem partitioning, familiar to anyone who is adept at troubleshooting (or anyone who has played "animal,

vegetable, mineral"), lets you manage large numbers of rules, while still honing in on the problem quickly.

Expert systems generally require a human to build and structure rules. These systems excel when there are many well-defined rules that cover most or all possibilities. However, they aren't good at responding to unanticipated situations, as the rules won't necessarily apply in those cases.

Neural networks, on the other hand, rely on nodes called *perceptrons*. The perceptrons are designed to simulate the actions of human neurons. Basically, a perceptron generates an output when specific input conditions, or thresholds, are met. By itself, a perceptron can only solve certain classes of problems. However, networks of perceptrons can effectively solve lots of real world problems like optical character recognition and credit card fraud detection.

For instance, a neural network for detecting credit card fraud might examine the card holder's balance, the average charge, and the holder's home zip code alongside the details of the current transaction. The network then assigns different weights to the various inputs and uses

them to form an opinion about the legitimate nature of the charge.

If the network is incorrect (that is, it marks a fraudulent transaction as legitimate or vice versa) the system can be trained to correct its mistake. This is the advantage of the neural approach—networks can learn. Learning takes place when perceptron thresholds are adjusted by algorithms that compare the network's output state to the desired state.

The disadvantage of such a network is that, like people, networks take awhile to learn and occasionally make mistakes. Algorithmic training may take many iterations to converge on a set of thresholds that work in most or all cases. Also, just as in real life, faulty training produces faulty networks, and separating good training from bad can be difficult.

Thus, while products from the likes of HP, IBM, Lightflow, and others claim to have AI, don't forget that AI technology covers a lot of ground. If you want to separate the facts from the hype, you'll need to dig a little deeper. Don't be afraid to ask the vendor just what technologies it's actually employing and why it chose them.

—Al Williams

For students at Harvard Medical School, staying on top of the course material is hard enough without having to track information at 50 different locations. Students stay organized and focused via a wireless network and specially formatted content for their mobile devices.

case study

Joseph C. Panettieri

a prescription for wireless

The Harvard Medical School system keeps students current

Harvard Medical School is known for its rigorous academics, but that isn't the only challenge for students at the prestigious university. Until recently, students had a difficult time keeping track of class changes, schedule updates, pending tests, and other dynamic information from the school.

At any one time, the students occupy 50 different locations as they rotate between their daily classes, hospitals, libraries, and living quarters. John Halamka, associate dean of Harvard Medical School and CIO of CareGroup, a healthcare organization that serves six Boston, MA-area hospitals, explains that most students used reams of paper to repeatedly print and update daily agendas, school announcements, and surveys from the university. "There was a clear need for a mobile solution," says Halamka.

The medical school wanted to build an application that let mobile students access course calendars, announcements, class notes, and syllabi. In addition, students wanted to log patient observations into their devices. Although Halamka knew the medical school couldn't permit the handheld systems to access patient records due to privacy and legal concerns, the other features they needed to add already amounted to a big project.

evaluation

In February 2001, Harvard Medical School began evaluating possible solutions to help students stay organized and connected while on the go. Among the questions that

Halamka and his team had to decide on were whether they should create custom, text-based content for Palm handhelds, and whether they would have to write custom applications in C for Palm-compatible devices. They also considered adopting a Microsoft .Net infrastructure for mobile systems, as well as repurposing content from the Web and somehow delivering it to cross-platform handheld systems.

Each of the approaches had potential benefits and drawbacks. For instance, if they had created custom content for Palm handhelds, the user experience would have been very rich, but that approach had two major pitfalls. First, the Palm-native content would have to be reconfigured for use on laptops and Windows CE-based handhelds. Second, it would have been difficult for Harvard Medical School to maintain and dynamically update separate content environments for both the Web and handhelds.

Programming C applications for the Palm using tools from Metrowerks (now owned by Motorola) had similar drawbacks. "Programming in C was the nastiest option," says Halamka. "It would have taken forever, and it would have been a Palm-centric approach." Again, the difference between Palm devices, laptops, and CE-based handhelds would pose a problem. With this route, the group would have had to rewrite their applications for each device.

Harvard's third option—embracing Microsoft's .Net strategy—also prompted skepticism at the medical school. On the

upside, Microsoft's Mobile Information Server (MIS) lets developers extend corporate Windows applications to mobile wireless devices. And the price was reasonable for the group: MIS costs \$15 to \$20 per user and requires a nominal dual-process Pentium II server running Windows 2000.

The catch? Because of its Microsoft roots, MIS was designed to work best with Windows CE-based handhelds. Microsoft is quick to note that the system also supports third-party browsers and cell phones. However, Harvard didn't think that MIS offered tight integration with Palm-based handhelds.

MIS has other potential drawbacks for cross-platform customers, as well. For instance, the application requires customers to activate at least one Microsoft Active Directory server within an organization. Although many Microsoft shops are embracing Active Directory, the software is far less mature than the likes of Novell Directory Services (NDS, also known as eDirectory). NDS has shipped with NetWare for nearly ten years and is now available for Linux, Solaris, and Windows servers.

Active Directory, by contrast, debuted with Windows 2000—roughly two years ago. Many customers are still trying to master Microsoft's directory service. According to a recent IDC survey of 300 IT managers, about 36 percent of Microsoft customers have delayed Windows 2000 deployments because of the complexity associated with Active Directory.



Next Stop
New York

Digital Video
DV
expo
EAST

June 25-26, 2002
Jacob Javits Center
New York, NY

DVexpo.com/east

Image from Artbeats Digital Library

In conjunction with

pc
xpo

TECHXNY



CMP

United Business Media

Digital Video
DV
Media Group

Priority Code: NAR

web repurposing

Harvard ultimately decided to pursue its fourth option, a Web repurposing approach that formatted existing content for delivery to handhelds and other mobile systems. The solution was built with software from AvantGo and wireless LAN hardware from Cisco.

"AvantGo is the industry leader in Web-to-mobile synchronization," says Halamka. "I had used them in the past to synchronize *The Wall Street Journal*, *The Boston Globe*, and *The New York Times* to my personal Palm."

Although AvantGo's technology appears to be sound, the company's financial performance could be cause for concern among some customers. AvantGo lost \$56.2 million during 2001. However, the company says it has enough funding to operate for six more quarters, and Harvard has no plans to abandon the vendor's mobile software.

AvantGo's flagship product is M-Business Server, a software suite that runs on Linux, Sun Solaris, and Windows NT/2000 servers. The software supports key wireless protocols such as CDMA, CDPD, GPRS, GSM, and 802.11b. It also works with an array of wireless devices—including Palm and PocketPC-based handhelds from Casio, Compaq, Hewlett-Packard, IBM, Sony, and Symbol Technologies, among others.

Harvard Medical School decided to outsource the integration work to ArcStream Solutions, a consulting firm in Watertown, MA. "ArcStream really understands the user interface issues on the mobile devices," says Halamka about the decision.

Although ArcStream is only two years old, the company's management team has strong relationships with Harvard University and many CIOs throughout New England. ArcStream CEO John Keane graduated from Harvard Business School in 1988. Moreover, John Keane Sr. serves on the board of CareGroup, where Halamka is CIO. Palm Chief Internet Officer Barry Cottle sits on ArcStream's board.

The work began in early 2001 as collaboration between ArcStream and Harvard's IT team. ArcStream contributed one project manager and one primary developer, while Harvard assisted in the effort with two of its own developers.

The project's biggest hurdle, according to Halamka, was finding an easy way to synchronize the handheld systems with

the back-end applications. An obvious choice was to use docking stations for the handhelds, but that wasn't an ideal solution because each device vendor makes a proprietary docking station for its respective device. Instead of using the docking stations, the medical school embraced the standardized infrared synchronization capabilities that are built into all major handhelds.

wireless infrastructure

In addition to the IR synchronization for handhelds, Harvard wanted to set up a wireless LAN so that students with laptops could access data from any location. Harvard evaluated hardware from Cisco and Nortel Networks, and ultimately chose Cisco's Aironet 350 product line because of its "more robust feature set," recalls Halamka. The Aironet family has found a home on numerous campuses, including the University of Akron, the University of North Carolina, and Western Michigan University.

Like most wireless LAN technology on the market today, Aironet adheres to the 802.11b standard overseen by the IEEE. Such products are designed to communicate wirelessly at 11Mbps.

However, according to a study by technology distributor Ingram Micro, the actual throughput is typically closer to 4Mbps.

Harvard deployed eight Aironet 350 wireless LAN access points in its five-floor medical center and two-floor library. Most distributors sell the access points for roughly \$939 per node and \$200 per PC adapter card. According to Halamka, the access points and adapter cards work "quite well" as long as they're within 500 feet of one another. He also notes that the Harvard network typically delivers a throughput of 6Mbps, and to date, dropped connections haven't been an issue at the school.

ensuring security

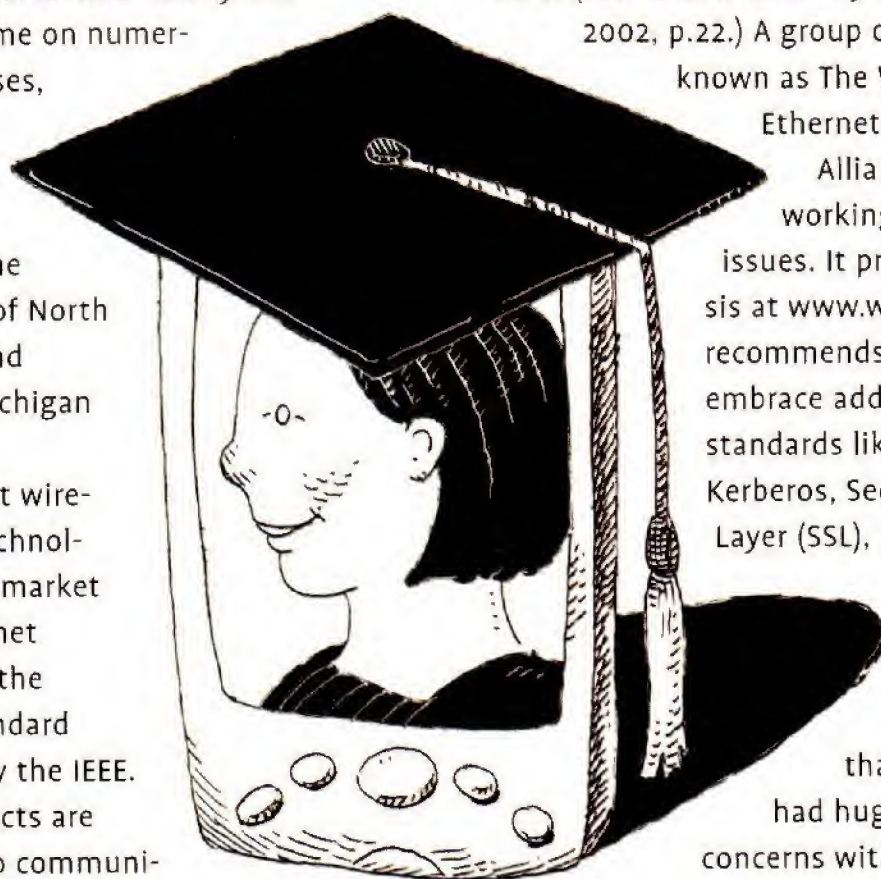
Like many potential customers, Harvard had security concerns about the wireless LAN infrastructure before the purchase. Indeed, many IT managers fail to activate or properly configure security features when they deploy wireless LANs. In an example that was recounted in numerous security-related newsletters earlier this year, James Taschek, a technology consultant in San Francisco, accessed a bank's wireless LAN from a bus stop situated across the street from the bank. Taschek's hack of the bank was accidental. He says he was working on his laptop when his system's wireless LAN card suddenly "discovered" the bank's wireless network.

Meanwhile, customers and consultants alike have uncovered several other security issues with the 802.11b wireless LAN standard. (See "LAN Sharks" by Paul Sholtz, May 2002, p.22.) A group of vendors

known as The Wireless Ethernet Compatibility Alliance (WECA) is working to address the issues. It provides an analysis at www.wi-fi.org. WECA recommends that customers embrace additional security standards like Radius, Kerberos, Secure Socket Layer (SSL), and firewalls.

Harvard Medical School concurs with that advice. "We had huge security concerns with this project," recalls Halamka. "But we took the

appropriate steps to address them." The school uses standards like IPSec for protected tunneling and authentication across the wireless network. Harvard also deployed security hardware from Bluesocket, a wireless LAN specialist in Burlington, MA. Bluesocket's hardware is similar to a firewall that sits between a company's wireless access points and its core, wired network. If a hacker breaks into the wireless system, Bluesocket's hardware prevents the hacker from probing wired networks and connected back-end systems. Halamka says he is pleased with the security solution that his group has implemented.



Aside from problems with security, wireless LANs often suffer from interference issues. Extensive metal construction, like elevator shafts, often blocks or degrades wireless LAN signals. Microwave ovens are also known culprits.

Halamka says that Harvard Medical School has suffered no such interference problems, but concedes that the wireless system and related PDAs aren't without their problems. "You have to recognize that wireless handheld technology is still bleeding-edge stuff," he says. "Every week that goes by it gets more and more stable and robust. But just a year ago, software drivers and endless code changes could create major headaches in a wireless system. This area of technology is evolving so rapidly that you have to tread warily."

a successful rollout

After implementing a pilot program that began on May 4, 2001 and ran until August 29, 2001, Halamka's team introduced the MyCourses system to Harvard Medical School. Today, more than 700 students, armed with laptops and handhelds, use

Harvard Medical School's wireless network and its related mobile applications.

Students are pleased about the time they save with MyCourses. "I definitely find that MyCourses...maximizes my time as a medical student," says Jeffrey Chung, a member of the class of 2004. "It allows me to access important course material daily while I'm in transit to the hospital or classes; and I especially enjoy not having to input course contact information myself," he adds. "The application saves me a lot of time."

Most students use the infrared connection to synchronize their handhelds with data on Harvard Medical School's servers. But more and more handhelds will likely embrace the 802.11b wireless LAN standard for higher-speed connections. Wireless LAN cards are now coming to market for PocketPC handhelds. Halamka expects similar cards for Palm-based devices to emerge later this year.

The school may also upgrade from 802.11b to newer 802.11a wireless LAN gear, which works at an advertised 54Mbps. A spokesperson for Proxim, a wireless LAN

vendor in Sunnyvale, CA, says that 802.11a throughput is roughly 20Mbps, or more than four times that of 802.11b. However, most 802.11a and 802.11b gear cannot be mixed and matched on the same network due to compatibility issues.

Nevertheless, Harvard Medical School may embrace 802.11a to push video streaming out to laptops. "One hour after every class ends, the associated lecture is available as a video stream at Harvard," notes Halamka.

In the meantime, the \$250,000 wireless handheld system continues to pay healthy dividends. In its first year of use, the system saved \$150,000 in paper costs. "Several schools have asked us how we did it," says Halamka. "We're not selling the system, but I'm more than happy to offer advice to those who are seeking a similar solution." ■

Joseph (joe_pan5@yahoo.com) is a ten year veteran of business and technology journalism and editorial director at the New York Institute of Technology.



Advanced Hosting Platform

- Windows 2000 Server
- SQL Server 2000
- MS Access 2000
- ASP / XML / WAP
- ASP Components
- Live Statistics
- Active Directory
- Windows Media
- HostMechanic™

Cutting-edge Publishing Tools

- Visual Studio.NET
- MS FrontPage 2002
- Microsoft® Office XP

Unleash the Power



Visual Studio .NET

iNNErHOST makes it easy to manage your .NET web applications by fully supporting the integrated publishing capabilities of Visual Studio .NET. Look for the iNNErHOST logo in your Visual Studio .NET installation to get started, TODAY!

- .NET Framework
- .NET Mobile Internet Toolkit

- XML Web Services
- ASP.NET Components

Now offering a FREE Visual Studio .NET hosting account for a limited time!

Win a FREE copy of Visual Studio .NET

Visit: www.innerhost.com/winvs

Copyright © 2001 iNNErHOST Inc. All rights reserved.






Toll Free: 888.751.5274

www.innerhost.com

Complex Hosting in a Global Environment™



Bret is an intellectual property and Internet attorney with Hancock, Rothert & Bunshoft. You can reach him at bret@lertext.com.

legal
code

Bret A.
Fausett



captured The sad, predictable story of ICANN's at-large membership

The Internet Corporation for Assigned Names and Numbers (ICANN) was created in 1998 as part of the Clinton administration's effort to privatize management of Internet resources and "lessen the burdens of government." It may have seemed an odd move for a Democratic administration to move oversight of a quasi-public resource from the government to the private sector, but neither the administration nor those on the outside pushing for privatization were blind to the public interest implications.

That the public had a vested stake in the sound and stable operation of the Internet's core functionality wasn't even open to debate. The original ICANN bylaws called for the election of eighteen board members, nine from industry and nine from a proposed at-large membership that would be representative of the Internet's end users. No one knew exactly how the at-large Internet community would be organized or represented on the board, or how members would be selected. But everyone agreed that half of ICANN's board of directors would come from the worldwide community of Internet users. This promise was at the very heart of the Commerce Department's agreement to accept the ICANN proposal to manage the Internet's domain name system.

An initial board of nine directors, appointed by the Internet's longtime Domain Name and IP Address Administrator Jon Postel, was charged with getting ICANN off the ground and creating the mechanisms for seating a fully elected and accountable board. That initial board started two

independent tracks to develop ways to elect the nine Board members from industry and nine Board members from the Internet community at large.

No one knew it in 1998, but this two-track strategy would eventually kill the concept of the at-large in its entirety... and maybe ICANN with it.

building the at-large

One of ICANN's first official acts was to create a Membership Advisory Committee (MAC) to consult the community and advise the board on how to create ICANN's at-large membership structure. This committee confronted significant problems. Who would be allowed to join? How would they join? Would there be a minimum age? How would individual members be authenticated? How would those who joined as members elect their designated seats on the board of directors?

Perhaps the most vexing question facing the MAC, however, was how to prevent the board from being captured by a determined minority. The power to elect half the board was the power to rule it. If we let the world's Internet users vote on all nine

seats, would the United States—with the most wired population in the world—dominate the voting, and hence, the seats elected by that voting? At some point in the future would we see an ICANN Board with nine at-large members all from China? What would prevent an interested corporation from enlisting its employees to enroll as voters in an attempt to gain a board seat, or multiple seats, for corporate management?

All of these were difficult questions without obvious answers, and it was important to get the answers right the first time. As with most corporations, two-thirds of the ICANN Board could change the corporate bylaws at any time. And nine board members elected by an aberrant election process would be close to the number required to capture the board and change the bylaws in their favor forever.

When the MAC presented its final report, it recommended that elections take place within regional voting pools and that voting be open to any individual who wished to participate. Corporations and organizations wouldn't be eligible to vote

The nine directors elected from the supporting organizations didn't share the public interest values inherent in the initial conceptions of ICANN, and their actions have shown that.

Visit **new.architect** magazine online!



Featuring:

- Complete editorial content
- Daily blogs from the editors and industry experts
- IT white papers and market research reports from over 60 top analyst firms and 3,500 IT vendors
- Industry event listings
- Job and resume postings
- Complete Web Techniques archives
- And much more!

Sign up for:

- A **FREE** subscription to *New Architect* magazine
- *This Month at NewArchitectmag.com*, our monthly email newsletter containing article summaries, links to online resources, updates on current events, and exclusive content you won't get anywhere else

newarchitectmag.com

as corporations or organizations, but individual employees would be welcome to enroll and vote as individuals. The MAC recommended that the most effective means of ensuring that a determined minority didn't capture the election process was to enroll a significant number of voters. It recommended outreach to the Internet community as the means to that end.

building the supporting organizations

The MAC was scheduled to finish its recommendation to the board in time for the at-large elections to take place within ICANN's first year of life. The committee finished its work on time, in just six months.

By the time the MAC had completed its report, however, the initial board of directors had recognized three supporting organizations corresponding to ICANN's core areas of responsibility: domain names, IP addresses, and Internet protocols. Each supporting organization would be primarily responsible for developing policy within its defined subject matter, and each would elect three directors to sit on the ICANN Board.

Given the MAC's recommendation to prevent capture of the at-large by enrolling a significant number of voters, the nine initial ICANN directors decided to proceed with outreach to the Internet community before electing the at-large board members. Several of the early participants in ICANN's activities doubted that anyone cared enough about domain names and IP addresses to vote in an election of directors. Because of this concern, outreach was designed to ensure that when ICANN finally held its at-large elections, a sufficient number of voters actually showed up to vote.

Although outreach would delay election of the at-large board members, the nine supporting organizations directors still would be elected in time for ICANN's annual meeting in November 1999.

Although everything seemed to be in order, things changed after the supporting organizations elections. The nine directors from the industries regulated by ICANN now balanced the initial nine directors appointed by Postel. And so all eighteen were responsible for how—and whether—the at-large membership was implemented.

why have an at-large?

One of the directors elected by the Protocol Supporting Organization was Dr. Vint Cerf, the man often described as the "father of the Internet." From the beginning, Cerf was dubious of the need for an at-large membership within ICANN, and his views were typical of those elected from the Supporting Organizations side of ICANN. He once likened ICANN to a public utilities commission or state transportation department: "The idea that a particular infrastructure service, kind of like the power system or the road system, needs to have a kind of global election to manage the underlying infrastructure is hard to fully appreciate. We don't do that for the power system, and we don't do it for the road system."

He also believed that the great unwashed might be incapable of electing knowledgeable directors, and this feeling was also pervasive within the newly constituted board of eighteen. The board initially thought it could solve the problem of an ignorant electorate by creating a blue-ribbon nominating committee that would pick the candidates. Only after opposition from public interest groups, like Common Cause and the Center for Democracy and Technology, did the board relent and let the Internet community add its own nominations to those of a nominating committee.

The board also decided that the first elections from the at-large membership would only elect five Directors. The Supporting Organizations had elected nine of their members, but the at-large would now have to proceed in baby steps.

the first elections

The initial belief that ICANN's activities were so arcane that no one would be interested in participating in its online election was challenged in the first days after registration opened. Thousands of people signed up. When the ICANN election became well known, tens of thousands more signed up. Registration campaigns in Europe and Asia even created apparent competition among nations to see which country could register the most voters. By the time the registration period ended, over 150,000 people had signed up. These numbers dwarfed even the most optimistic expectations about interest in ICANN.

Half of the people who registered to vote took the affirmative steps necessary to confirm their identities and establish actual eligibility to vote, and half as many again actually voted in the election. Still, having over 35,000 people participate in an online, worldwide election was a stunning success. As a direct result of those elections, five board members took seats in November 2000.

The new board of ICANN was composed of nine directors from the supporting organizations, five from the at-large community, and four holdovers from the initial Postel appointments.

Had the first elections been a success? Three of the elected directors came from the slates of the ICANN Board's nominating committee, but the two that hadn't weren't exactly what the existing directors had hoped for. North America elected Karl Auerbach, a long-time ICANN activist and dissident. Europe elected Andy Müller-Maguhn, a leader in Berlin's fabled Chaos Computer Club. The technical credentials of both Mr. Auerbach and Mr. Müller-Maguhn were impeccable, but their points of view were far from those of the other directors. That shouldn't have been a surprise, of course. The views of the directorate now more accurately reflected the disparate views of the Internet's users.

studying the options

Depending on how one viewed the outcome, the first elections were pronounced a success, a failure, or just an interesting experiment. Because there was no agreement on how to proceed and how to elect the final four at-large members, the board commissioned another blue-ribbon committee to study what had happened with the first election and plot a way forward. This group, called the At Large Study Committee (ALSC), was headed by Carl Bildt, former Swedish Prime Minister and UN envoy to the Balkans. Its panel of participants was impressive.

With a half million-dollar budget and the assistance of a full-time staff person, the ALSC started what it referred to as a "clean sheet" study of the at-large. Every aspect of the at-large would be re-examined.

After months of deliberation, however, the ALSC members unanimously agreed that the at-large community of Internet

users was an important stakeholder, and deserved both meaningful participation in ICANN's processes and representation on ICANN's board.

How to accomplish those goals was the hard question. The ALSC ultimately recommended a compromise in which one-third of the board would come from the at-large community, with another third coming from suppliers of DNS services, and the final third from developers of DNS services.

But even that compromise proved too difficult for the ICANN Board.

the palace coup

After the ALSC had submitted its final report—but before its recommendations could be acted upon by the Board—ICANN's President Stuart Lynn presented what he called his "Case for Reform." This was a plan to dismantle all of ICANN's structures and rebuild them anew. Lynn's plan called for the elimination of the at-large membership. While ICANN would create new avenues for "participation," it would abandon the idea that representation on the board was attendant to that. The official public voice would come from giving the

governments of the world a say in board member nominations, but online voting would never occur again. Lynn suggested replacing the at-large directors with a self-perpetuating board nominated by a board-approved committee.

Those who had worked hard to make the first elections a success were clearly dismayed by the Lynn reform plan. Perhaps no one expressed that sentiment better than ALSC member Charles Costello, director of the Carter Center's Democracy Program. He named the plan a "palace coup d'état" and argued that it was "a breach of faith with the founding principles and basic structure of ICANN as well as the fiduciary duty of [the] board...[to] so fundamentally change the governance structure of the organization and permanently disenfranchise an at-large membership." Meeting in March 2002, half a world away from ICANN's home in California, some board members feigned insult at the charge that they were coup plotters. However, it didn't change their course of action. They moved forward on the Lynn reform plan, calling for the creation of participatory structures, but not the creation of new election structures.

On one level, it certainly looks like a palace coup d'état. But in retrospect, the outcome of these last three years of at-large development efforts was entirely predictable. Three years ago, the MAC contended that the most important aspect in the creation of an at-large was that it shouldn't be susceptible to capture. While it was going about its work, however, the ICANN Board itself was captured by the early elections from the supporting organizations. The nine directors elected from the supporting organizations didn't share the public interest values inherent in the initial conceptions of ICANN, and their actions have shown that.

ICANN may have killed the at-large membership at its March 2002 meetings in Ghana, but it was diagnosed as terminal over two years ago. ICANN was built to provide private-sector management of private resources, but with the public interest hard-wired into its deliberative structures. Now that the initial vision of ICANN has died, one has to wonder whether what remains will live for very much longer either. An effective voice for the public interest was always a critical element of ICANN's mission. ■

71 Issues of *Web Techniques* Magazine at Your Finger Tips (No shelf-space required)



You may not have kept your back issues of *Web Techniques*, but *New Architect* has. Search the complete archive of issues from April 1996 to February 2002. Don't miss this essential online resource for Web strategists, authors, designers, and developers.

Visit the *Web Techniques* magazine archives today:
newarchitectmag.com/archives/



CMP

United Business Media

To avoid hiccups, don't go live before putting your systems through the necessary load tests. Using a dedicated hardware-based solution for load testing not only gives you more control than software solutions, but you'll also get a more realistic mix of traffic and stress.

review
::hardware

see how it feels to be popular Caw Networks WebAvalanche 4.0

What's the worst way to load test your Web servers? Put your new system live on the Internet, and then wait for it to become popular. That way, when you make the big time, your server and site fail. Of course, you can avoid such a fate by investing in a load testing system.

Caw Networks's WebAvalanche 4.0 (I tried the 1800 model) is a comprehensive hardware solution—a complete tester-in-a-box. I was impressed by it, but it comes with a hefty price tag (\$19,500). So, I also looked at some software-only solutions as alternatives (see the “Software-based Load Testers” sidebar).

WebAvalanche is powered by a dual 1GHz Pentium III machine with five 10/100 interfaces. Four are dedicated to testing, and one is for remote management. A version with a 1Gb fiber interface replacing the four 10/100 cards is also available. The system runs the QNX real-time operating system, and a custom IP stack controls the Ethernet interfaces. Software upgrades are handled by replacing a CompactFlash card. An internal hard drive is provided for storing test results.

Setting the system up was easy. You plug it in, hook up network cables, and power on. WebAvalanche is a dedicated lab test system, and is intended to be used on an isolated network. Plug the four test interfaces into a standalone Ethernet switch and put the servers you want to stress test onto the same switch. Then, plug the management Ethernet interface into a port on your LAN. (I had to adjust the IP addresses on my subnet a bit to access the WebAvalanche interface.)

WebAvalanche runs a Web server to allow control of all features. It's possible to put a monitor and keyboard on the WebAvalanche, but you wouldn't normally need either because you'll probably end up using the Web interface most of the time. The Web interface lets you set up and store profiles for virtual users and test scenarios. It provides a fine-grained level of control over tests. And you can view progress via real-time displays in Java applet windows.

power to burn

Caw Networks claims that WebAvalanche can simulate 1 to 2 million simultaneous users. To use as a target for the tests, Caw loaned me a 1.5 GHz P4 server running Red Hat Linux. I ran Apache on it, both serving a small HTML page and running a simple Perl CGI script. I wanted to test the network interfaces, so I hit one URL repeatedly in each test sequence. Apache serves the hits out of memory, so server disk I/O becomes less of a factor in testing.

One approach to using a stress tester is to hit the server hard enough to make it start generating errors, then tune the server to improve its performance. Setting a test

to ramp from zero errors to overload means that you can methodically tweak server settings and then re-run the test to see if you've made any improvement.

Ramping up from 10,000 to 20,000 users (generating around 1,500 to 2,000 transactions per second) was enough to cause a 50 percent or more error rate. As you might guess, hitting the URL for the CGI script was much harder on the server. It supported only about 200 transactions per second before errors started to outnumber correct responses.

Reports from WebAvalanche are available either as graphs generated while the test is underway, or for entire test results, as CSV files that you can save to your desktop and/or view with Excel. Caw also provides an optional accessory program, called Icepick, for viewing the files.

WebAvalanche pretends to be a whole Internet worth of traffic. You can configure ranges of IP numbers to originate traffic and control data rates for each range to simulate different line speeds from slow modems up to maximum wire speed. Caw offers a companion product, called WebReflector, that pretends to be a whole slew of Web servers. You put other network components between the two, such as routers or firewalls, to see how they hold up.

The user interface on WebAvalanche could use more work. For example, you have to click a few times to re-run the same test. Caw plans to add a Copy Test button that will let you duplicate an existing test, and then modify its settings instead of starting from scratch.

At a Glance

WebAvalanche 4.0

Company	Caw Networks
URL	www.caw.com
Price	Starts at \$19,500, and options with streaming start at \$39,500.

WebAvalanche is impressive. (But it certainly isn't the only load tool available—like WebAvalanche, all of the software products listed in the sidebar to this article can playback lists of URLs and handle basic authentication, cookies, and Web forms.)

WebAvalanche can also test QuickTime media streams—up to 10,000 concurrent RTSP/RTP streams. For power, flexibility, and ease of set up, WebAvalanche wins hands down over similar software solutions. Straight out of the box, I could generate the

heavy loads required to test and overload a pretty beefy server. No other product I tested comes close. ■

—Brian Wilson

Brian is cofounder of Harbro Systems. Write to him at bwilson@harbrosystems.com.

software-based load testers

Although hardware-based testing devices can be convenient, software packages offer many of the same features for less.

empirix e-test

The e-Test suite of programs includes e-Test for verifying site functionality and e-Load, the load tester. To run a load test, you run e-Test as a proxy server, and it records the URLs as you access them from a browser. You can monitor a running test in graphs. As with the other software-only products I tested, the trick is to generate enough load from a desktop machine. You can use e-Test as a control console, and run multiple instances of the e-Load Agent on several workstations, too.

The e-Test suite also includes Web site monitoring and report generator modules. Empirix tried to cover a lot of ground with this product. The regression-testing ability is e-Test's strength.

mercury interactive

Of the software-based testers, Mercury Interactive's Astra LoadTest stands out clearly as the best. LoadTest was easy to install, and the basic tutorial is well written. As with e-Test, LoadTest uses a browser proxy to let you specify a test data set. After saving the test data, you specify the rest of the test settings. Every imaginable parameter is controllable.

While designing a test, you can run in Debug mode to watch pages fly by in the browser window, then switch to Load mode when you're ready to run a load test.

LoadTest is completely Microsoft-oriented. You can export data only in Excel format, and the Load Generator can only run on Windows systems. The HTML files it generates use backslashes, so they can't be viewed on non-Windows machines without editing them.

web performance trainer

I tried the 2.3 version on Windows 2000 and Linux. As with e-Test, Trainer runs as a Web proxy server. To enter the information for a test, you configure a Web browser to use Trainer as a proxy, tell Trainer to record, and then browse the URLs you want added. This is called "creating a business case." After you have set up one, you can incorporate it with other cases, or directly play it back.

Also like e-Test, Trainer can be configured to run in an agent mode, where one station controls several systems running the playback engine. Trainer agents can run on a mix of platforms. For example, you could run Windows on the control station and Linux on headless agent systems.

Setting up and running tests with Trainer is fine, but its reporting system is poor. All other commercial systems provide a set of canned graphs. With Trainer, you must set up the parameters you want displayed. The resulting graphs are unattractive so I suggest you use another graphing tool to process results exported from Trainer in CSV files.

webserver stress tool

The Paessler Webserver Stress Tool is simple to set up. A navigation bar on the left of the screen selects test set up, user settings, and test results panes. All configuration settings for a given test set are stored together in one scenario file.

Webserver Stress Tool supports the most significant features of the more expensive products. Its biggest limitation is that it can't control multiple agents running on other systems, so the load you can generate is limited to the speed of the desktop machine on which you run it. It does let you use multiple Ethernet interfaces.

Setting up a test works as it does with the other products, by recording browser activity, or you can type in URLs directly. As test runs are executed, test results scroll by in text format.

Results are shown in both text logs and graphs, and can be directly exported to a Word document or printed. Unlike the other products here, Webserver Stress Tool doesn't generate CSV files.

siege

The truly budget conscious can go the open source route. I tried out Siege on Linux. It's a command-line tester, so setting up, modifying, and repeating tests is easy, and it's free. With a little extra scripting, you should be able to run it on multiple servers to generate heavier traffic loads. Though results are in text form, Siege lets you export its data to CSV files for further processing.

—BW

Load Testing Software at a Glance

	Astra LoadTest 5.4	e-Test Suite 6.0	Webserver Stress Tool 5.30	Web Performance Trainer 2.3	Siege
Company	Mercury Interactive	Empirix	Paessler	Web Performance Inc.	open source
URL	www.merc-int.com	www.empirix.com	www.paessler.com	webperformanceinc.com	freshmeat.net/products/siege/
Price	\$29,995 for 250 virtual users Windows NT/2000	\$5,000	\$325	Minimum \$1,495 for 100 virtual users.	n/a

HTML's great advantage was never its slick navigational interfaces for Web applications, but authoring simplicity. Flash MX may be the answer for rich application front ends.

review ::software

build rich front ends to your web applications Macromedia Flash MX

When Macromedia acquired a little vector-based animation program called FutureSplash in 1997, it renamed the program Flash. At the time, few imagined that Flash would eclipse all other plug-in content technologies, even Macromedia's own, once-flagship product, Shockwave. Flash has always been noted for its small player size and tiny content downloads, which let you deliver sophisticated presentations to even modem users. In Flash 4—released in 1999—Macromedia added a scripting language, ActionScript, that allowed for rudimentary control of Flash movies.

But as the Web morphed from delivering static pages (or static Flash movies) to dynamic applications, developers and designers found that Web browsers had limited user interface functions. The Web's request and response model is inherently ill-suited for real-time data updating, because HTML pages have to reload in full every time a browser checks for new data on the server. As we move to the world of Web applications, it becomes clear that HTML isn't always the best tool for providing a robust front end for real-time Web applications. Enter Flash MX.

rich media gets richer

From the multimedia perspective, the big news is the new video support. Thanks to Sorenson Spark, a new codec from the makers of the one used in QuickTime, you can import AVI, DV, MPEG, QuickTime, and Windows Media File (Windows only) movies. (For more details, see the sidebar "Video Support with Sorenson Spark.") When you select the Import Video command, you see the Sorenson Spark dialog box, which lets you set the compres-

sion quality, keyframe rate, and scaling. You can also set the video to synchronize with the frame rate of your Flash presentation. (Don't do this if you want to be able to drop frames.) Another switch lets you choose whether to input the audio track.

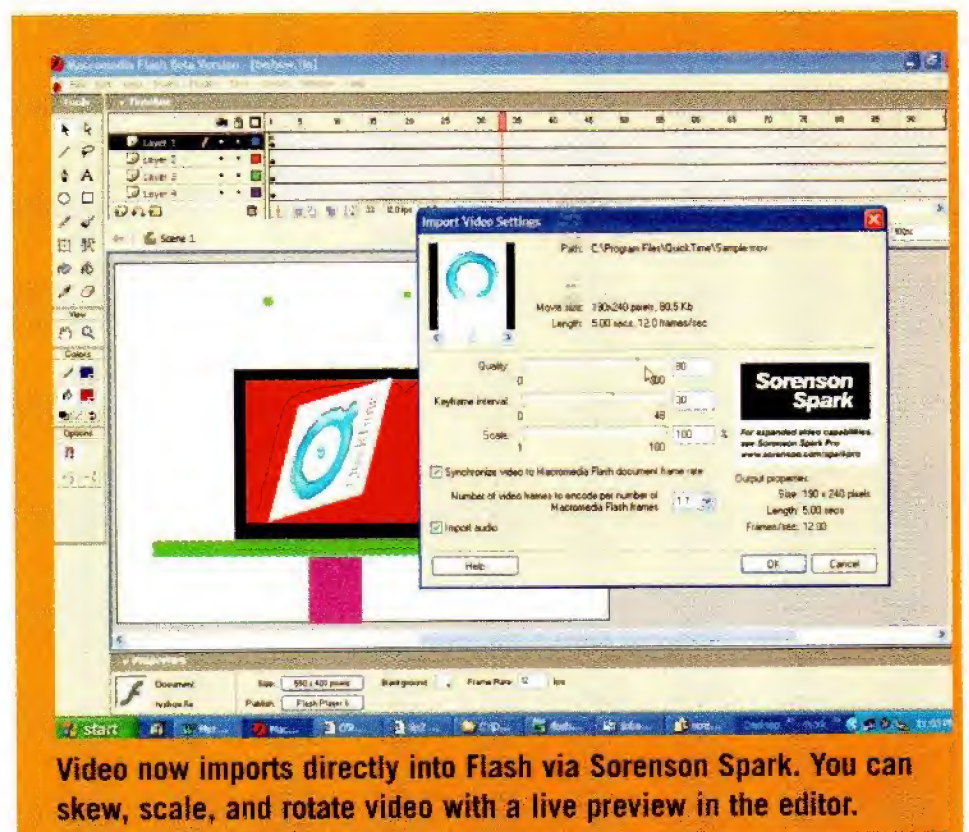
Once video is imported, Flash treats it like any other imported object. In a quick test, I imported the 81K QuickTime sample movie (as in the figure), skewed and rotated it, and exported it as a Flash .swf movie. I then reduced the file size to 24K using a relatively low quality setting. The compression factor (better than 3:1 in that case) is quite substantial. Impressively, the video plays at full speed in its transformed state in the editor. Video can also be sourced in at runtime via ActionScript, so your movie doesn't have to balloon in size to have video. (Indeed this is a consistent theme in the new version. Not only video, but JPEGs, sounds, and other ActionScripts can be sourced in at runtime.) If you combine video with the various animation effects for which Flash is so famous, Flash starts to look like a post-production facility on a PC. Of course, you can export Flash as QuickTime or AVI.

There are other useful tool enhance-

ments. The text tool now features vertical type and fine kerning control. A snap-to-pixel feature lets you align graphics perfectly. Flash MX also has a Break Apart function that converts type into shapes, and lets you manipulate individual letters. In addition, Flash now supports dynamic text, so a live Flash file can read in a text (or XML) file from the server during runtime. There are also several UI improvements like panel docking, better color management, a new Properties panel, and the ability to nest Timeline layers in folders.

components

The other major breakthrough is the Components panel. In MX, forms elements are little self-contained Flash movies that you can drop onto the main authoring window (the Stage). Built-in components



Video now imports directly into Flash via Sorenson Spark. You can skew, scale, and rotate video with a live preview in the editor.

Include check boxes, radio buttons, scroll bars, list boxes, and drop-down lists. Rather than coding all of the parameters for a given object, you can just fill out the Properties panel for a selected component. To connect a form to a server-side script, you can write ActionScripts to communicate via HTTP, just like a Web browser. With shared UI components, Macromedia hopes that Flash apps will have consistent interfaces for users down the road.

Another interesting facet of components is that their appearance (skins) can be changed. If you create new graphics for the components, all instances of the component in your Flash document will display your new skin. Finally, via ActionScript, you can store the form input in an object and validate user data before sending to a server. Beyond this, each component has an API that lets ActionScript handle the behavior programmatically.

actionscript

In the previous version of Flash, ActionScript was less than a full-featured scripting language. In Flash MX, ActionScript is a full-blooded scripting language based on ECMAScript, just like JavaScript. Every object in the Flash environment is now exposed through an API, so Flash developers are to some degree freed from the tyranny of the

GUI. For example, to change an object's position, color, and transparency, why go through the authoring process of setting keyframes in the timeline, visually moving the object, and changing its properties by hand, when a simple function in a script will handle it?

MX also exposes methods and properties of just about every element in a Flash movie to ActionScript. The ActionScript editor features some nice hinting tools, and a useful debugger.

better and faster

One of the litmus tests for clients these days is the ability to interact with Web apps. Thus, Flash's support for XML is crucial to its role as a serious front end. Flash offers an XML object for requesting data via HTTP's PUSH method, and an XMLSocket for maintaining open connections.

Macromedia is developing a Web application strategy that will include a Flash application server gateway, and new server components for ColdFusion, J2EE, and .Net. Also look for Macromedia to deliver a messaging format so that ActionScript can perform remote procedure calls (RPCs) with Web applications and servers.

With Flash MX's strong support for Web standards and XML, and its vector drawing mechanics, it's a little surprising that Flash

At a Glance

Flash MX	
Company	Macromedia
URL	www.macromedia.com
Price	\$499
Pros	Powerful new version of ActionScript; video support; improved ActionScript editor; performance gains; rich, new Components Panel feature.
Cons	No SVG support.

doesn't export to SVG, an XML standard for representing graphics. Flash MX would be a great SVG editor.

Taken together, the rich multimedia plus the increased scripting support and hooks into middleware and server products suggest not only that some very cool content is headed your way, but also that Flash is becoming a feature rich and robust client application development platform. ■

—Richard Koman

Richard is a writer and an editor. You can contact him at rkoman@sonic.net.

video support with sorenson spark

Flash MX comes with the built-in ability to import video into your Flash movie, thanks to the Sorenson Spark codec. Sorenson is responsible for the codec in QuickTime, so it knows what it's doing with video compression. Both Spark and Squeeze, Sorenson's advanced video compressor, are based on the new Sorenson codec: Sorenson Video 3 Pro. When you import a video (AVI, QuickTime, or DV) in Flash, a dialog box lets you set quality, number of keyframes, and scale. There are also options for matching the video frame count to the Flash movies, and importing audio. While the Squeeze in Flash does a fine job of compressing video—especially at quality levels over 70 percent—check out Sorenson's Squeeze product if you want more control over compression.

you can import it into Flash as you would with any other digital video file. Squeeze uses the Sorenson Video 3 codec for all video compression, and a choice of four audio codecs. For audio, you can select the data rate, bit rate, and whether the sound is mono or stereo. For video, you can specify the data rate, the frame rate, the encoding method (constant bit rate, one or two-pass variable bit rate), as well as a number of more advanced options—such as hinting for streaming, bidirectional prediction of interframes, masking, and more. Bidirectional interframes (B-frames), basically the ability to take any two keyframes and calculate the interframes in either direction, are key to Sorenson's ability to deliver good compression without necessarily reducing the frame rate. In addition, you can tell Squeeze to add hints to video for

Squeeze is a \$299 tool that allows you to convert AVI, QuickTime (.mov), and DV files to QuickTime, Flash movie (.swf), or Flash video (.flv) formats. Once you have a QuickTime or .flv file,

streaming. Sorenson is also launching a streaming hosting service called Vcast, which wasn't available as of this writing.

You start Squeeze by opening a video file (AVI, QuickTime, or DV). Next, you choose a compression option from the menu. Squeeze comes with presets for modem, ISDN, broadband (low, medium, and high), LAN, and CD. The Edit menu lets you change these default values to create your own compression settings. For example, you might want to alter the data rate or the frames per second on the video side, or reduce the sampling rate on the audio side. Once your options are all set, you simply choose your setting and click a button.

For comparison, here are some of the file sizes created from the different default settings:

- Modem: 289K,
- ISDN: 870K,
- Broadband-Low: 1.8MB,
- Broadband-High: 3.7MB,
- CD: 10.6MB.

—RK

WAP-based application development hasn't gone away. Two new toolkits offer polished features that you'd expect from a standard SDK. You only have to choose your target device.

review
::software

hope for wireless developers

WAP toolkits from Nokia and Openwave

The last three years or so have been a rather giddy ride for WAP developers. Although it has been a messy process, WAP has not only survived, but morphed into a set of rich, useful technologies. Still, support in the form of SDKs and dedicated authoring tools has been frustratingly lacking. As WAP enters a new stage in its development, however, the various tools that are available have followed suit with update after update.

Two of the most popular development kits are Openwave's Software Development Kit and the Nokia Mobile Internet Toolkit 3.0. Both are freely available to developers, and between them they support the majority of the WAP-enabled handsets sold today.

openwave sdk 5.1 preview

Openwave offers two SDKs for download (versions 5.0 and 5.1). The SDK version 5.0 helps you create WML 1.3-compliant applications (but with Openwave extensions), plus WMLScript 1.2 services for the Openwave Mobile Browser WAP Edition 5.0. The SDK version 5.1 Preview helps you create XHTML Mobile Profile 1.0 services with Cascading Style Sheets (CSSs) for the Openwave Mobile Browser Universal Edition. I reviewed version 5.1.

The SDK 5.1 Preview illustrates just how much the GSM industry is riding on the success of the GSMA M-Services Initiative, an industry commitment to mobile architecture. A lot of thought has clearly gone into this excellent, intuitive software.

The SDK includes several features to aid developers in targeting their applications towards Openwave's proprietary browser, called Mobile Browser:

- a text editor with syntax highlighting that supports XHTML, WML, and WMLScript content creation;
- a debugger, and testing tools with trace output (supported by Siemens S45 and the Generic 5.0 device);
- various device simulators with embedded Mobile Browser;
- Openwave Mobile Access Gateway server simulator;
- support for Download Fun (Openwave's implementation of the GSMA M-Services Initiative);
- all of the usual sample code, help files, and supporting documentation.

The SDK installation was straightforward, but the setup program required a restart, which was mildly annoying. You may have to change your firewall settings to give the SDK access to the Internet.

Upon launching the SDK, your default browser opens and connects to an Openwave site that has demo content. The examples are simple, and won't be informative for anyone but a novice user. The integrated editor is useful for minor work and the syntax coloring is always a welcome addition, but as editors go, it's limited. There are no project, source control, or workflow features—all key elements of enterprise packages. It

depends on your personal preference, but I think many users may find an external editor easier to use when constructing WML documents, and the SDK editor better for syntax checking and device emulation.

The SDK really comes into its own (as an SDK should) when you need to resolve errors. You can watch the handshake output between device and server, view errors, and trace the output. You can view source for hints about the composition. You can also monitor cookies and history to get your applications working smoothly.

Switching between various SDKs and their programming tool of choice may be annoying for some. It's especially difficult for people who are used to creating dynamic content with Microsoft's Active Server Pages for .Net. Visual Studio .Net is an astounding application and is pretty much *de rigueur* for creating ASP.Net applications. However, Microsoft's vision of mobile services—and thus, the tools it offers for their creation—differs considerably from that of Openwave. It remains to be seen whether the two will compliment or hinder each other.

A major benefit of the SDK is that, subject to free developer registration through the Openwave developer Web site, you can access a live provisioning gateway—very useful for testing.

As a standalone application, the SDK is competent, and the developer resources available through Openwave's developer relationship program are extensive.

nokia mobile internet toolkit 3.0

At first, the Nokia Toolkit was disappointing. This slow-loading Java application has separate windows scattered around with confusing contents. And, why a package that only runs on Microsoft platforms needs to be written in Java is a mystery to me, given the hit on performance. In fact, waiting one second for a menu selection to take effect rapidly became frustrating. I found that altering the Look and Feel (in the Preferences menu) from Java to Windows speeded things up a bit. There were no such difficulties when I used the Openwave SDK.

That aside, many options that build on previous releases of the toolkit are available to the developer. These include:

- various device simulators;
- WML and WMLScript encoders;
- WML Decks, WML Scripts, WBMP Images, Multipart Content, XHTML, and CSS editors;
- WAP Protocol Stack, HTTP, and file access modules;
- debugging views;
- Nokia SoftID, a WIM card simulation for developing security applications;
- example applications and source code;
- and documentation and context-sensitive help.

The toolkit installed without a problem and didn't require a restart, which earned it good marks right away.

Not all views within the toolkit are available at all times, as these are dependent on the capabilities of the device itself. For example, the content, history, and active variable weren't available in the 3330 emulator. The comprehensive content view pane contains the resulting decoded WML, the bytecode, an element tree that recreates the document structure, and facilities for viewing the original source file.

The built-in Gateway emulation operates a serviceable environment to test WML applications. However, you can also use a direct HTTP connection, or specify a proxy. The toolkit includes a WAP emulator that's compliant with the June 2000 release, and a Nokia Mobile Browser simulator. You can download additional device emulators from Forum Nokia. Emulators that are currently available include Nokia 3350/3395, 6210, and 7110 devices, with more to follow. These provide true emulations of the devices and are invaluable aids to the developer.

Panels offer educational views of real-time compilation errors; the current document source (whatever the content type); context sensitive history; variables in use and their values (which can be changed to observe the results); session views where HTTP handshakes can be monitored; and a toolkit log that displays processing information at a level that's determined by the user. These panels are invaluable for tracking down and resolving sometimes complex code problems.

All in all, the toolkit is essential for improving your understanding of the way different Nokia devices will handle your

application. However, the toolkit has some rough edges and a build quality that doesn't give me much confidence in its robustness for everyday use.

different objectives

Openwave is a stronger SDK, but the two products have different objectives. Each attempts to aid the creation of WAP applications specific to Openwave's and Nokia's very different technologies. Thus, developers seeking to create cross-platform applications—especially those devices that aren't covered by Nokia and Openwave browsers—might do well to seek help elsewhere. Several familiar applications will help create WML documents, including Macromedia's Homesite and Dreamweaver, and Adobe's GoLive. But at the enterprise, vendor-neutral level, there really isn't much choice. The situation deepens when you need to delve into WAP's more exotic standards.

Worse, should Microsoft start making inroads into the mobile market, developers will be faced with learning Microsoft's often impenetrable .Net strategy. Add in other hardware vendors like Alcatel, Mitsubishi, Motorola, NEC, Sagem, Siemens, and others, and it quickly becomes apparent that we need an application that will meet the needs of all these devices. As the mobile market continues to diversify, the merits of producing SDKs and toolkits with a tunnel-vision attitude toward creating applications can only serve to turn people away.

Which toolkit you'll choose will depend on the business factors involved, such as whether you're dealing with a service that will host your content on Openwave gateways, or targeting the youth sector that may be brandishing shiny new Nokias. Both SDKs are quite capable of making what could be a painful experience much easier to deal with. In the context of SDKs, there's a Catch-22 between the real world and the ideal world. The Openwave SDK and Nokia Toolkit both take an admirable stab at aiding beleaguered WAP developers. ▀

—Richard Weeks

Richard operates BrightFluid, a UK research consultancy that studies behavioral patterns behind mobile communications usage. You can contact him at richard.weeks@brightfluid.com.

Two WAP Toolkits At a Glance


	Openwave SDK 5.1 Preview	Nokia Mobile Internet Toolkit 3.0
Company	Openwave	Nokia
URL	www.openwave.com	www.forum.nokia.com
Price	free	free
Pros	Access to an Openwave gateway for live application testing. Well designed, intuitive IDE, and excellent supporting documentation.	Built-in gateway. New phones are available as plugins.
Cons	Limited to handsets containing the Openwave Mobile Browser. WML editor isn't of much use.	Java-based IDE can be slow. Limited to Nokia handsets. Messaging functionality not directly supported, so you'll need to download another toolkit.

New Architect asked User Experience Architect Alex Wright to review a favorite book from his bookshelf that has influenced his work over the past year. He was inspired by this book about IDEO, a company that has successfully and consistently excelled at the creative process.

review
::book

inspirational guidance *The Art of Innovation: Lessons in Creativity from IDEO*

How has IDEO sustained such a consistently creative, productive, and profitable culture over the past twenty years?



Tom Kelley, general manager of IDEO (and brother of founder David Kelley), and his co-author Jonathan Littman answer that question in *The Art of Innovation: Lessons in Creativity from IDEO*, America's Leading Design Firm (2001, Doubleday & Company, \$27.50).

When AT&T hired Palo Alto-based IDEO to develop new concepts for its teleconferencing service, the IDEO team started out by doing what many user-centered design teams might do: it talked to users.

But not just any users. Instead of hunting for "typical users"—the pre-screened, carefully segmented, demographically consistent humdrums who tend to fill up focus groups—IDEO sought out the "crazy user."

For example, an office worker named Sally had developed her own idiosyncratic approach to teleconferencing that involved carting a bunch of individual speakerphones into a conference room, setting the phones around a table, dialing in each caller separately, and conducting the conference call in the open air of the room.

Eccentric? Sure. But Sally's unorthodox behavior yielded breakthrough insights for the design team. Her need to control the physical dimensions of the conference room sparked the IDEO team's creativity. As a result, the company developed a new

teleconferencing product that incorporated numerous physical and spatial cues into the product design.

That kind of unique thinking has propelled IDEO to near-mythic status in the industrial design world. The progenitor of such varied products as the Apple mouse, Crest toothpaste caps, the Palm V, the Heartstream heart defibrillator, and the Aerobie football, among thousands of others, IDEO has mastered the art—and the business—of innovation.

The Art of Innovation, a chatty manifesto about creativity and organizational culture, overflows with anecdotes about the many design challenges IDEO has faced over the years. For instance, the one about how an early prototype for the Apple mouse took shape out of a butter dish and a rolling ball. Or the one about how a team working on a ski goggle project in Palo Alto in mid-July faced the challenge of mimicking winter fog conditions. The solution: place a bicycle in an ice cream freezer, and recruit willing colleagues to don goggles and pedal until they break a sweat.

Kelley and Littman bring readers inside the IDEO culture by organizing the book around discreet themes, each supported by vignettes from the IDEO project archives.

As you might expect of a ghost-written encomium penned by (and for) a company executive, *The Art of Innovation* reads at times like a book-length commercial for IDEO. Readers in search of company gossip or juicy tales of what it's really like at IDEO will, of course, be disappointed. But despite the book's relentlessly positive viewpoint, it serves up plenty of good advice on how to imbue any organization with the creative spirit.

The authors also talk about the importance of what they call "cross-dressers," or team members who switch disciplines or specialties. Engineers turned designers, for instance, and vice versa. IDEO blurs disciplinary boundaries wherever possible. That's especially true when it comes to research, a cornerstone of IDEO's design process. Many design firms still treat research as a stand-alone discipline practiced by Researchers (with a capital R). At IDEO, research is everyone's job.

However, the book's most concrete lessons are in the chapter on brainstorming. Here, the authors expound on a series of simple ground rules for brainstorming. For example, an optimal brainstorm should last one hour and should have a specific goal—say, of coming up with one hundred ideas (physically numbered on a big sheet of paper). Brainstorms should never be scheduled off-site, lest creativity be construed as some kind of unusual act that could only ever happen outside the office.

Toward the end of the book, the authors stray from the subject they know best—the IDEO way and its experiences with customers—into a mode of management soothsaying that, alas, devolves into predictable platitudes and proscriptions better left to the likes of "trend expert" Faith Popcorn. But when the authors stick with the core topic, the story of IDEO, this book brims with great accounts and practical insights that far outweigh the occasional predictable lapses into company boosterism. ■

—Alex Wright

Alex is a user experience architect. You can contact him at alex@agwright.com.

Less Than \$10*

Internet orders only, shipping and handling included

MSDN Magazine

2001 CD-ROM

**All 2001
Editorial
Content, with
Source Code,
on One CD!**

also includes Complete article
index for *MSDN Magazine*,
MSJ and *Microsoft Internet
Developer* from 1986 to 2001

Order Today!

msdn.microsoft.com/msdnmag/cdrom/2001



* MSDN Magazine 2001 CD-ROM is \$9.95 for all Internet orders placed by July 15, 2002 including shipping and handling for orders delivered within the U.S. only. Extra charges for multiple copy orders, rapid delivery, and delivery outside the U.S. will apply; exact charges will appear when online order is placed. CD-ROM orders placed by phone (1-800-444-4881) or mail (MSDN Magazine CD-ROM 2001, 1601 West 23rd Street # 200, Lawrence, KS, 66046) are \$25.95 including shipping & handling. Mail orders accepted for U.S. delivery only and must be accompanied by a check for \$25.95. Credit cards accepted on Internet and phone orders only.

Lincoln is an M.D. and Ph.D. who designs information systems for the human genome project at Cold Spring Harbor Laboratory in New York, NY. You can contact him at lstein@cshl.org.

crow's nest

Lincoln D. Stein



the morpheus incident How corporate squabbles could stifle the Web

At the end of February, the peer-to-peer world was rocked by a bizarre series of events. First, a million users of the Morpheus file-sharing application suddenly found themselves locked out of the network. After initially blaming the problem on a technical glitch, the company that makes Morpheus, StreamCast of Tennessee, claimed that it had been sabotaged by Australia-based Sharman Networks, proprietor of the FastTrack file-sharing protocol. StreamCast then issued a "preview edition" of its version 2.0 product, which on closer inspection turned out to be a rebranded version of Gnutella, an established Gnutella client.

Meanwhile, Kazaa, a competing file-swapping service based in Holland, began a "Welcome Over to Our Place" campaign directed at Morpheus users. It posted a software installer that would "fix" Morpheus users' software by migrating them to the Kazaa system.

What's going on? The answer is ugly. After the Recording Industry Association of America closed Napster down, several companies jumped in to claim Napster's place in the music and file-sharing niche. Eventually, two camps emerged: products built on the open-source Gnutella P2P protocol, and products based on the proprietary FastTrack protocol. FastTrack services include Kazaa (the original inventor of the FastTrack protocol), StreamCast's Morpheus, and another service called Grokster.

Because StreamCast licensed its technology from a competitor in the music-sharing industry, there may always have been some tension between it and Kazaa. However, the situation

became worse in December 2001, when Kazaa, threatened by a copyright lawsuit in the Netherlands, folded and sold its name and assets to Sharman Networks of Australia.

Sharman seems to have been less enthusiastic about sharing its technology with StreamCast. According to StreamCast's CEO, when the FastTrack protocol was updated to version 1.5, Sharman Networks failed to provide StreamCast with the information it needed to take advantage of the protocol's new features. StreamCast protested this action by withholding some \$60,000 in licensing fees.

The dispute simmered until the end of February, when StreamCast claims that Sharman Networks pulled the plug on Morpheus without warning. Using a facility in the FastTrack protocol that was designed to trigger live software updates, Sharman propagated a message across FastTrack that caused Morpheus clients to try to update themselves. But because there was no new version of Morpheus to update to, the software was effectively disabled.

This action left StreamCast scrambling to pick up the pieces. Unable to patch things up with Sharman, StreamCast turned to a ready-made alternative, the open source Gnutella client. StreamCast pasted its logo onto the Gnutella splash screen and released the software as a preview product. This was allowed under the Gnutella GNU General Public License (GPL) terms, provided that the modified source was also redistributed. Although initially neglectful of this obligation, StreamCast released the modified source code to its Web site a few days later.

Whether StreamCast will be able to retain its user base is a matter of speculation, but let's think for a moment about the implications of the Morpheus incident. Due to a technical dispute between two companies, a million users permanently lost the use of a piece of software. It's like AT&T retaliating against Southwestern Bell by listing a fuse that melts all of the telephone handsets made by Southwestern Bell. Under the license agreements that software companies use, this type of action is legal. In fact, the proposed Uniform Computer Information Transactions Act (UCITA) encourages the use of software time bombs to enforce license terms.

As the Web moves toward a world of distributed services, companies will increasingly have both the motive and opportunity to interfere with competitors. Consider strategic choke points like Microsoft's Passport service or AOL/TimeWarner's cable routers. If Microsoft can make competitors' Web sites seem slow to accept credit cards, or AOL/TimeWarner can cut off competitors' streaming video, how long do you think they will resist the temptation to do so?

There are solutions. Government regulation could force companies to keep the Internet open. Or we could defend and expand the use of open source protocols on the Web to prevent any one company from gaining a choke hold on the network through proprietary technologies.

The Morpheus incident is a wake-up call. We must maintain the Internet as a level playing field where all comers have an equal opportunity to compete. Only in this way can new technologies like P2P grow and flourish. ■

ad index

Advertisers in this issue

ADVERTISER	URL	PAGE
Adaptive Path	www.adaptivepath.com	27
Advanced Internet Technologies	www.aitcom.net	23
Apple Computer	www.apple.com	C2-1
Computer Associates	www.ca.com/etrust	4
Covalent	www.covalent.net/ers	11
Innerhost	www.innerhost.com	41
Interland	www.interland.net	C4
Intermedia	www.intermedia.net	6
Microsoft	msdn.microsoft.com/vstudio	2-3

This index is provided as a service to our readers. *New Architect* assumes no liability for errors or omissions.

Advertising contact information

Sales Director

James Long
415-947-6421
jlong@cmp.com

Eastern Region Sales Manager

Chris Salem
203-798-2215
csalem@cmp.com

Western Region Sales Manager

Michelle Stutsman
415-947-6416
mstutsman@cmp.com

Southern Region Sales Manager


James Long
415-947-6421
jlong@cmp.com

Ad Production Coordinator

Paul O'Hare
415-947-6941
pohare@cmp.com

Publisher

Bob Melk
415-947-6425
bmelk@cmp.com



New Architect magazine
CMP Media LLC
600 Harrison Street
San Francisco, CA 94107
Tel: 415-947-6000
Fax: 415-947-6026

newarchitectmag.com/advertising/

Steven is CTO of hesketh.com/inc in Raleigh, NC. He is also the author of *Building Dynamic HTML GUIs*, and has been technical editor for several other books. You can write to him at schampeo@hesketh.com.

guest
editorial

Steven
Champeon



why dhtml will win

Competition among user interface tools heats up

There's a lot of bold talk coming from a certain multimedia tools vendor (Macromedia, cough, cough) lately, about how its new Flash MX product is "the future of the Internet." Never mind that the company leaders seem to be confusing the Internet with the Web. What's interesting is how they demo this rich, multimedia future. The vendor's Web site makes much of an ETrade stock quote application—something that could have been thrown together in half an hour with Dynamic HTML (DHTML) without the need for proprietary technology, plugins, or a massive press campaign. How very 1997.

DHTML, you say? Tried that, hated it. It didn't work the way you had expected in [insert your favorite browser here], right? Well, it's time to take a second look. What you may think of as DHTML has advanced greatly since the early days of Web design. Back then, Microsoft and Netscape were in a pitched battle over which Document Object Model (DOM) would rule the Web: `document.all` or `document.layers`. But that has all changed. In the end, reason and the W3C have prevailed. Even Netscape, the company that created the `<layer>` element, has abandoned it in favor of the standard `<div>`. Developers no longer have to code for two different APIs.

So why the resistance to DHTML? Perhaps it's time for a new name. The past six years have brought enormous change, standardization, and stability (not to mention top-notch implementations) from the major browser vendors. And that isn't just because "Microsoft won the war." We all won the browser wars. Compromises were made on

both sides, and the Web of tomorrow will bear as much resemblance to the antebellum days of Mosaic as your desktop resembles that of Bartleby the scrivener.

Because DHTML uses the components that your browser already supports, like HTML, Cascading Style Sheets (CSS), JavaScript, and the DOM, it requires no fancy authoring environments or server side applications. You don't need to learn another programming language or scripting variant, as some client GUI technologies require. You don't need to worry about the distinction between source files and binary files. And the sites that you produce can have a clean separation between document structure, presentation logic, and behavior—all of which let teams focus on their strengths, instead of forcing designers to program, or programmers to design.

What's more, the DOM and CSS were designed to work well with XML, a true enabling technology for which you should be preparing your company. (The latest version of HTML, XHTML, is already an XML document type. Thus, it's clear where things are headed.) Just try sharing data like that with a Flash file, or linking to a specific page in a Java applet.

DHTML solves business problems. It can be used to augment traditional, document-oriented sites as well as serve as the platform for a new breed of Web-based applications.

When was the last time you needed scalable graphics and typography, animation, or sound to solve your business problems? (Unless you're in the online porn industry, or the Web multimedia authoring tools business, of course.)

You might object to this perspective on the grounds that we do, in fact, need client-side processing, personalization, reliable client/server communication, and a way to save state without disrupting the page display. But DHTML lets you do all of that. And it doesn't require your customers to sit through needless delay while your plugins or virtual machines load.

Another of DHTML's advantages over its competitors is that you needn't worry about locking yourself into a proprietary platform. If you're building an application or updating your site, it's your decision whether to use a commercial, disruptive multimedia format owned by a single company and ruled by its financial ideology, or to adopt an open, standardized platform.

If you have a serious problem to solve, then the question is not "How should the Web evolve?" Rather, it's "How has the Web already evolved, and how can I take advantage of it now?" So, come back to Dynamic HTML for the very first time. You may be surprised at just how much DHTML resembles the Web we have right now, and vice versa. ▀

When was the last time you needed scalable graphics and typography, animation, or sound to solve your business problems?



Our Content Management System Buyers Kit won't make you a superhero.

Just a brilliant decision maker in office appropriate attire.

Choosing the right content management system is complex, but you don't need special powers. Just get the Content Management Buyers Kit. It's based on the experiences of people who have gone through the evaluation process, and emerged victorious. Its tools will help you build your evaluation team, ask the right questions, assess your company's needs, and understand the various solutions. So you, too, can make the right decision. And if that's not enough to make you look like a hero, the kit is free.

Get the free Content Management Buyers Kit.
www.atomz.com/nw



You think one dedicated service manager is revolutionary? Try two.



Baron M.
*Dedicated
Account Manager*

Thomas D.
*Dedicated
Support Manager*

Get the most service and support from Interland's Reseller Program.

• NEW, exclusive reseller support

Our customer service is about you. Your clients. Your business. We promise that you'll experience superior customer service. With more dedicated support managers focused on you, you'll receive the highest level of support possible. Getting unparalleled support will keep your clients happy and your business growing.

• Dedicated account and support managers

Not one but two dedicated managers will be there to help you at a moment's notice. Your account manager will focus on your overall business, making sure your clients are happy and your accounts are working properly. Your support manager will ensure technical problems won't come between you and your clients. A personal pager number or cell phone number guarantees they're within reach. And each manager is backed by a team of the industry's top professionals.

• Put more money in your pockets

Reselling Interland hosting plans is a profitable business. With our award-winning products and services, you can increase your profits with each additional hosting plan you sell. The more you sell, the deeper the discounts. The deeper the discounts, the more money in your pocket.



TRUEADVANTAGE™

Shared	Number of plans	Discount	1-Year Service Agreement	2-Year Service Agreement
Preferred	(4-10)	20%	25%	28%
Select	(11-24)	25%	30%	33%
Premium	(25-39)	30%	35%	38%
Ultimate	(40+)	35%	40%	43%

Designed to help your customers build an online presence, TrueAdvantage hosting plans provide painless setup and include easy-to-use control panels.

FREEDOM™

VPS	Number of plans	Discount	1-Year Service Agreement	2-Year Service Agreement
Preferred	(4-10)	30%	35%	38%
Select	(11-24)	40%	45%	48%
Premium	(25-39)	45%	50%	53%
Ultimate	(40+)	50%	55%	58%

Built to give you the control you need, Interland's Freedom hosting plans empower you with industry-leading Virtual Private Server technology, root access and multiple domain hosting.

Visit interland.com or call us at 1-866-846-7789.

*Offer valid through August 15, 2002. Offer is valid only on new Power 200 and Power 300 plans with a 12-month or longer account. Offer is valid on all new TrueAdvantage and Freedom Plans with a 24-month pre-paid account. Free months are added at end of full contract account period. Setup fees will apply. Promotional offer valid only for telephone sales. Promotions cannot be combined with any other offers. © 2002 Interland, Inc. All rights reserved. Not responsible for errors or omissions in typography or photography. Interland is a registered trademark and the Interland design, Power, PowerView, Web Extras, True, TrueAdvantage, TrueBasic, TrueAdvanced, TrueBusiness, Freedom Accelerator, ezSiteStarter and ezSiteBuilder are trademarks of Interland, Inc. Microsoft is a registered trademark of Microsoft Corporation. All other product names appearing herein are for identification purposes only and may be trademarks of their respective companies. All orders subject to credit approval, availability, and other terms and conditions (including terms of service and acceptable use policy) available at www.interland.com. Features, functions and pricing subject to change without notice.

**GET 3 FREE
MONTHS OF
WEB HOSTING**
for each new Interland account you open.*